



LANCE MENTHE, DAHLIA ANNE GOLDFELD, ABBIE TINGSTAD, SHERRILL LINGEL,
EDWARD GEIST, DONALD BRUNK, AMANDA WICKER, SARAH SOLIMAN, BALYS GINTAUTAS,
ANNE STICKELLS, AMADO CORDOVA

Technology Innovation and the Future of Air Force Intelligence Analysis

Volume 2, Findings and Recommendations



For more information on this publication, visit www.rand.org/t/RR341-2

Library of Congress Cataloging-in-Publication Data is available for this publication.

ISBN: 978-1-9774-0633-0

Published by the RAND Corporation, Santa Monica, Calif.

© Copyright 2021 RAND Corporation

RAND® is a registered trademark.

Cover: U.S. Marine Corps photo by Cpl. William Chockey; faraktinov, Adobe Stock.

Limited Print and Electronic Distribution Rights

This document and trademark(s) contained herein are protected by law. This representation of RAND intellectual property is provided for noncommercial use only. Unauthorized posting of this publication online is prohibited. Permission is given to duplicate this document for personal use only, as long as it is unaltered and complete. Permission is required from RAND to reproduce, or reuse in another form, any of its research documents for commercial use. For information on reprint and linking permissions, please visit www.rand.org/pubs/permissions.

The RAND Corporation is a research organization that develops solutions to public policy challenges to help make communities throughout the world safer and more secure, healthier and more prosperous. RAND is nonprofit, nonpartisan, and committed to the public interest.

RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

Support RAND

Make a tax-deductible charitable contribution at
www.rand.org/giving/contribute

www.rand.org

Preface

The Air Force Distributed Common Ground System (AF DCGS) is responsible for producing and distributing actionable intelligence from data collected by a wide variety of U.S. Air Force platforms for warfighters around the world. Over the past two decades, intelligence collections and the demand for intelligence products have grown exponentially, straining analytic capacity. At the same time, intelligence analysts often are too busy performing routine processing, exploitation, and dissemination (PED) tasks to focus on larger strategic analyses that may be required to meet future threats envisioned by the 2018 National Defense Strategy. A 2012 RAND Project AIR FORCE (PAF) report suggested that artificial intelligence (AI) would one day be able to help free analysts to do tasks that make better use of human intelligence. Since that report was published, AI and machine learning (ML) have made enormous advances, and we foresee further innovation in the coming years.

In 2017, Air Force/A2 asked PAF to analyze how current and potential future technologies could help AF DCGS become more effective, efficient, adept at using human capital, and agile. We were also asked to consider the process, training, and organizational improvements needed to make best use of these technologies. The research project, called *Closing the PED Gap*, was conducted in fiscal year 2018 in PAF's Force Modernization and Employment Program. The research is discussed in three companion reports:

- *Technology Innovation and the Future of Air Force Intelligence Analysis: Volume 1, Findings and Recommendations*, RR-A341-1, 2021. Volume 1 provides essential findings and recommendations for a broad audience, including Air Force decisionmakers.
- *Technology Innovation and the Future of Air Force Intelligence Analysis: Volume 2, Technical Analysis and Supporting Material*, RR-A341-2, 2021 (this report). Volume 2 provides more in-depth discussion of project methodology; a primer on AI and ML; more-detailed discussion of key recommendations; and other information that will be of interest to specialists, stakeholders, and experts.
- *Technology Innovation and the Future of Air Force Intelligence Analysis: Volume 3, Technical Assessment of Data Flow Maps*, forthcoming, Not available to the general public. Volume 3 furnishes additional restricted detail.

This report should be of interest to policymakers within the Air Force and the wider intelligence community.

The research reported here was commissioned by U.S. Air Force/A2 and conducted within the Force Modernization and Employment Program of RAND Project AIR FORCE as part of a fiscal year 2018 project *Closing the PED Gap*.

RAND Project AIR FORCE

RAND Project AIR FORCE (PAF), a division of the RAND Corporation, is the Department of the Air Force's (DAF's) federally funded research and development center for studies and analyses, supporting both the United States Air Force and the United States Space Force. PAF provides the DAF with independent analyses of policy alternatives affecting the development, employment, combat readiness, and support of current and future air, space, and cyber forces. Research is conducted in four programs: Strategy and Doctrine; Force Modernization and Employment; Manpower, Personnel, and Training; and Resource Management. The research reported here was prepared under contract FA7014-16-D-1000.

Additional information about PAF is available on our website:

www.rand.org/paf/

This report documents work originally shared with the DAF on October 3, 2018. The draft report, issued on September 26, 2018, was reviewed by formal peer reviewers and U.S. Air Force subject-matter experts.

Contents

Preface	iii
Figures and Tables.....	vii
Summary.....	viii
Acknowledgments	x
Abbreviations	xi
1. Introduction	1
Previous RAND Project AIR FORCE Work	1
Scope	3
Terminology	4
Research Methodology.....	6
About Data Flow Maps	8
2. Overview of the AF DCGS Today	13
Evolution of the AF DCGS and the Growth in Supply and Demand	13
Global Organization	16
Distributed Ground Stations Site.....	17
Crew Positions and Roles.....	20
Data Flow and AF DCGS Processes	22
AF DCGS Competencies	27
3. Improving Efficiency, Effectiveness, Human Capital, and Agility: Lessons from Historical Case Studies.....	31
Dynamic Targeting.....	31
Deliberate Targeting.....	36
Airman Resiliency	38
Humanitarian Assistance and Disaster Relief	40
4. Artificial Intelligence and Machine Learning: A Primer for AF DCGS Analysts	45
Historical Overview	46
Glossary.....	48
Current Status	50
Future Development	50
Artificial Intelligence/Machine Learning Methods	52
Supporting Technologies.....	59
Vulnerabilities	61
5. Improving GEOINT Analysis: Additional Detail	65
Making the Most of Today's Technology to Improve GEOINT Analysis	65
Taking Advantage of Future Technology to Improve GEOINT Analysis.....	70
Implementing GEOINT Recommendations	74
6. Rebalancing AF DCGS Competencies and Organization: Additional Detail	77

A Choice of Paths.....	77
Organizing Around Phases for Scalability	80
7. Building the Right Skills: Additional Detail	83
Building Programming and Data Science Skills	84
Sharing and Capturing Tactics, Techniques, and Procedures	86
Rehearsal-of-Concept Drills.....	86
Using Mission Type Orders and Focused Collections as Testbeds	88
8. Fostering Innovation and Successful Implementation: Additional Detail	89
Confronting the Innovation Adoption Problem	90
Stages of Technology Implementation.....	91
Existing Air Force and Department of Defense Efforts.....	100
Innovation Examples from the Intelligence Community	102
Conclusion.....	104
Appendix A. Defining Technology Readiness Levels for Artificial Intelligence/Machine	
Learning.....	105
References	109

Figures and Tables

Figures

Figure S.1. Major Recommendations	ix
Figure 1.1. Project Approach.....	7
Figure 1.2. Data Flow Map Symbols—Current Processes	9
Figure 1.3. Data Flow Map Symbols—Potential Future Entities and Connections	10
Figure 2.1. Major AF DCGS Sites	17
Figure 2.2. AF DCGS Operations Floor.....	19
Figure 2.3. Generic Data Flow Map—Current.....	23
Figure 2.4. DART Investigator Functions	25
Figure 2.5. Spectrum of Synthesis.....	27
Figure 2.6. AF DCGS Competencies	28
Figure 3.1. Find, Fix, Track, Target, Engage, and Assess Cycle	32
Figure 4.1. Layered Structure of Neural Networks	53
Figure 4.2. Recurrent Neural Networks	57
Figure 4.3. Tricking Machine Vision	62
Figure 5.1. Generic Data Flow Map—Potential Improvements with Today’s Technology	66
Figure 5.2. Generic Data Flow Map—Potential Improvements with Future Technology	72
Figure 6.1. Rebalancing Weight of Effort Across AF DCGS Competencies	78
Figure 6.2. Generic Data Flow Map—Potential Organization to Balance Competencies	81
Figure 8.1. Stages of Technology Implementation	91
Figure 8.2. Case Study: Pilot BlackBerry Use for Law Enforcement.....	95
Figure 8.3. Case Study: Health First	97
Figure 8.4. Case Study: Project Kessel Run.....	101

Tables

Table 1.1. Phases of Analysis.....	6
Table 2.1. Analysis Roles	21
Table 4.1. Glossary of AI/ML Terms	49
Table 5.1. GEOINT Recommendations: Making the Most of Today’s Technology	65
Table 5.2. GEOINT Recommendations: Taking Advantage of Future Technology.....	71
Table A.1. Proposed TRL Scale for AI/ML Methods	107

Summary

Issue

There is growing demand for the Air Force Distributed Common Ground System (AF DCGS) to analyze sensor data. Getting the right intelligence to the right people at the right time is increasingly difficult as the amount of data grows and timelines shrink. The need to exploit all collections limits the ability of analysts to address higher-level intelligence problems. Current tools and databases do not facilitate access to needed information.

Approach

Air Force/A2 asked RAND Project AIR FORCE (PAF) to analyze how new tools and technologies can help meet these demands, including how artificial intelligence (AI) and machine learning (ML) can be integrated into the analysis process. PAF assessed AF DCGS tools and processes, surveyed the state of the art in AI/ML methods, and examined best practices to encourage innovation and to incorporate new tools.

Conclusions

- Many analytic tasks can be fully or partially automated, although human involvement will continue to be necessary in more-complex tasks.
- AI/ML can free analysts to focus on solving intelligence problems and developing supporting technologies to make analysis more efficient.
- Analysts will require new skills both to facilitate use of AI/ML and to take advantage of opportunities to conduct more-advanced analysis.

Recommendations

AF DCGS should

- leverage existing technologies to automate some analysis and reporting tasks and to make archival intelligence more accessible
- take advantage of AI/ML technologies, when available, for early-phase analysis tasks (e.g., identifying and tagging imagery, issuing threat warnings, re-tasking collectors)
- organize to balance human effort across three competencies: supporting missions, supporting analysis, and solving intelligence problems
- recruit and train analysts with data science, programming, and other skills
- follow best practices for developing, implementing, and sustaining new tools.

Figure S.1. Major Recommendations

IMPROVEMENTS USING TODAY'S TECHNOLOGY		AI/ML-ENABLED IMPROVEMENTS	
GEOINT	Create a GEOINT Analysis and Reporting Tool to semi-automate product generation and mission reporting & assessment for EO/IR/SAR, FMV, MTI	GEOINT	Use AI/ML to perform partial first phase analysis (e.g. tag imagery, identify objects and people)
	Create Linker tool to tie information used to confirm the exploitation back to the source		Use AI/ML to alert human analysts of significant changes in activity
	Create improved Formatter for threat warning		Use AI/ML to generate threat warnings with human-on-the loop
	Create Updater script to automatically re-issue updated MTI threat warning messages		Use AI/ML to dynamically retask collectors
	Adopt Graphic Information System (GIS) into the MTI workflow and write Python scripts to automate GIS analysis processes		Use AI/ML to analyze archived FMV imagery
	Assess risks and benefits of adopting the industry standard in video editing tools (Avid)		Seek to lift the "eyes on" requirement for FMV missions where the "ISR role" indicates no risk of troops in contact or strike decisions, and when AI/ML tools can alert human analysts to other events that require real-time judgment
SIGINT	Create Scraper script to add warnings to reports	SIGINT	Leverage IC capabilities while maintaining organic capabilities for threat warning
	Build scripts and adjust workflow to transform or eliminate Technical Reporter position		
	Address barriers to using IC networks to pave the way for leveraging future capabilities		
Multi-INT	Reengineer and accelerate OA DCGS rollout	Multi-INT	Leverage IC capabilities for fusion and analysis
	Build on current efforts to link Crew Manning Letters, PED Tasking Orders, personnel qualifications databases, shift scheduling		Make cloud computing the centerpiece of the next hardware refresh

NON-MATERIEL IMPROVEMENTS			
Skills	Teach GIS basics at Goodfellow AFB		
	Encourage certificates for data science and programming, consider expanding Combat Readiness Sustainment Program, use time while analysts wait for clearances to expand skills and certifications		
	Retain basic INT skills for some Airmen, even where they may appear obsolete because of AI/ML		
	Support and expand rehearsal-of-concept drills		
	Hold an annual WEPTAC conference for sharing best practices among AF DCGS sites, particularly the DCGS Analysis and Reporting Team		
	Encourage use of mission type orders and focused collection operations		
Innovation	Emphasize user engagement at each stage; involve DGS-3 early in the development process of new tools		
	Identify opinion leaders and champions at DGS sites to foster tool development		
	Reward innovators with time to sustain their own innovations and to build new things		
	Create a process for gracefully offboarding old tools		

GEOINT=geospatial intelligence	SAR=synthetic aperture radar	IC=intelligence community
SIGINT=signals intelligence	FMV=full motion video	WEPTAC=weapons and tactics
EO=electro-optical	MTI=moving target indicator	DGS=Distributed Ground Station
IR=infrared	OA=open architecture	

NOTE: We recommend short-term improvements that can be implemented today and farther-term improvements that require technical breakthroughs or significant adjustments to AF DCGS network architecture and non-materiel improvements to make the AF DCGS more scalable while maintaining core capabilities. More detail on implementation is provided in Menthe et al., 2021; Menthe et al., forthcoming. The latter includes more-extensive discussion of SIGINT. AFB = Air Force Base; ISR = intelligence, surveillance, and reconnaissance; PED = processing, exploitation, and dissemination.

Acknowledgments

We would like to thank our sponsor, Michele Cook, Air Force/A2, for her guidance in shaping the project, and our action officers, Lt Col Brian Webster, Col Abraham Jackson, and Lt Col Steven Napier, for all their help along the way. We are also grateful to the leadership of the 480th Intelligence, Surveillance and Reconnaissance Wing (ISRW) and associated ISR Groups for opening their sites to us and allowing us to observe and speak to their analysts freely, sometimes for days on end. Thanks go to Col Jason M. Brown, Col David W. Dengler, Col Kristofer W. Gifford, Col Michael D. Holmes, Col Mark L. Mitchem, Col James C. Mock, and Lt Col Stacy S. Widauf. We also thank the leadership of 1st Special Operations Wing and 11th Special Operations Intelligence Squadron, Col Thomas B. Palenske and Lt Col Taylor O. Herron, respectively.

We are deeply appreciative of the time that so many analysts and other personnel afforded to us during our site visits, including: in Virginia, Capt Ryan Hess, Capt Jeremiah S. Johnson, TSgt Christopher R. Ewers, TSgt Joshua T. Price, SSgt Adam B. Carr, SrA Veach, SrA Ribyinski, A1C Brandon K. Ragoodial, and A1C Boris Warzecki; in California, Lt Sarah R. Farmer, TSgt Brandon Aguilar, TSgt Roy Davis, TSgt Lupe White, SSgt Ryan Gribble, SrA Jennifer Adams, SrA Pickett, and Richard Rivera, Jr.; in South Korea, Maj Timothy Cowan, Maj Stefan P. Katz, Maj Kalli J. Ritter, Capt Angela J. Ferreira, Capt Tranay Tanner, Lt Ian Heffron, SSgt Midnight D. Webber, and Alan Tokumoto; in Hawaii, Maj Jeffrey Dekkinga, Maj Chad Frey, Maj Joel Nieman, SMSgt Charles M. Edwards, MSgt Justin Gentile, TSgt Joshua Miller, TSgt Pyanowski, SSgt Paul D. Harper, Sgt Thomas R. Hileman, both Sgts Anderson, and SrA Alex J. Russell; in Indiana, CMSgt Thomas L. Bowker and MSgt Brent A. Evitts; in Georgia, Capt Cameron Shike, MSgt Michael Grorich, and TSgt Sara E. Davidson; and in Florida, Capt Joseph J. Mellone and Richard Custureri. Our apologies to the many others inadvertently omitted from this list.

We also want to thank several subject-matter experts we spoke to during the course of this work; specifically Dale Benedetti, Mark S. Brown, Maj Nicholas Schaeffer, Ryan Greer, and Kevin R. Carl at the National Air and Space Intelligence Center; Col Tim Gillespie at National Reconnaissance Office; and Nicholas A. Goralske and Manuel Arora.

Finally, we thank the many RAND colleagues who helped us with this work: Lt Col Matthew O. Berry for insights on special operations forces issues; Andrew J. Lohn and Gavin Hartnett for insights on AI/ML; and Elizabeth Bodine-Baron for insights on open-source intelligence. We also thank colleagues who answered our elicitation for organizational insights: Jair Aguirre, Michael Decker, Michael Hansen, Larry Hanser, Maria Lytell, Susan Resetar, Agnes Schaefer, Susan Straus, and Jennie Wenger.

There are far too many people for us to thank, including many who shared their insights with us in confidence. The efforts of all these people made this report possible.

Abbreviations

ACC	Air Combat Command
AF DCGS	Air Force Distributed Common Ground System
AFB	Air Force Base
AGI	artificial general intelligence
AI	artificial intelligence
ANG	Air National Guard
ART	Airman Resiliency Team
CAN	correlation analysis
CNN	convolutional neural network
COIN	counterinsurgency
CT	counterterrorism
CCMD	combatant command
DART	Distributed Common Ground System Analysis and Reporting Team
DCGS	Distributed Common Ground System
DGS	Distributed Ground Station
DIUx	Defense Innovation Unit Experimental
DMS	Distributed Mission Site
DoD	U.S. Department of Defense
EEI	essential element of information
ELINT	electronic intelligence
EMS	electronic intelligence mission supervisor
EO	electro-optical
F2T2EA	find, fix, track, target, engage, and assess
FMV	full-motion video
FY	fiscal year
GA	geospatial-intelligence analyst

GEOART	geospatial intelligence analysis and reporting tool
GEOINT	geospatial intelligence
GIS	geographic information system
GMS	ground mission supervisor
GRE	geospatial-intelligence reports editor
HA/DR	humanitarian assistance and disaster response
IC	intelligence community
IMS	imagery intelligence mission supervisor
INT	intelligence
IR	infrared
IRC	internet relay chat
ISR	intelligence, surveillance, and reconnaissance
ISRG	Intelligence, Surveillance, and Reconnaissance Group
ISRW	Intelligence, Surveillance, and Reconnaissance Wing
LSTM	long short-term memory
ML	machine learning
MSA	mission support analyst
MTI	moving target indicator
NASIC	National Air and Space Intelligence Center
NGA	National Geospatial-Intelligence Agency
NRO	National Reconnaissance Office
NSA	National Security Agency
OAF	Operation Allied Force
OEF	Operation Enduring Freedom
OIF	Operation Iraqi Freedom
OSINT	open-source intelligence
PAF	RAND Project AIR FORCE
PED	processing, exploitation, and dissemination
PMB	pre-mission brief

RNN	recurrent neural network
RoC	rehearsal of concept
RPA	remotely piloted aircraft
SAR	synthetic aperture radar
SIAS	sense, identify, attribute, and share
SIGINT	signals intelligence
SME	subject-matter expert
TEL	transporter erector launcher
TIC	troops in contact
TRL	technology readiness level
TSA	target system analysis
TTL	training task list
TTP	tactics, techniques, and procedures
UNICORN	Unified Collections Operations Reporting Network
WAMI	wide-area motion imagery
WEPTAC	weapons and tactics

1. Introduction

The U.S. Air Force/A2 asked RAND Project AIR FORCE (PAF) to investigate how technologies, tools, and processes can help the Air Force Distributed Common Ground System (AF DCGS) manage the growing demand for intelligence to support warfighters; make processes more effective, efficient, and agile; and make better use of human capital to meet evolving threats, such as those outlined in the 2018 National Defense Strategy. There are three volumes to this report; this is the second. Volume 1 presents findings and recommendations from this research.¹

This chapter provides additional detail about the research project. We first discuss how the project builds on previous PAF analyses of this problem. We then discuss the scope, define key terms, and describe the research methodology. Finally, we detail how we constructed the data flow maps referred to in subsequent chapters.

Previous RAND Project AIR FORCE Work

This research builds on several previous PAF projects stretching back almost a decade. PAF first examined AF DCGS operations in fiscal year (FY) 2009–2010 work on the intelligence, surveillance, and reconnaissance (ISR) force mix, in which processing, exploitation, and dissemination (PED) was identified as the bottleneck for counterinsurgency (COIN) and counterterrorism (CT) operations. In that work, PAF focused on full-motion video (FMV) operations and motion imagery processing and exploitation tools.² This was the first work in which PAF examined new tools and recommended organizing near real-time intelligence around geographic areas. In FY 2012, PAF looked at AF DCGS PED automation on a task-by-task basis and made recommendations for low-hanging fruit.³ In FY 2014, PAF examined how Air Force PED works within and makes use of the national signals intelligence (SIGINT) enterprise. PAF also looked at the special challenges of technical electronic intelligence (TechELINT). In FY

¹ Lance Menthe, Dahlia Anne Goldfeld, Abbie Tingstad, Sherrill Lingel, Edward Geist, Donald Brunk, Amanda Wicker, Sarah Soliman, Balys Gintautas, Anne Stickells, Amado Cordova, *Technology Innovation and the Future of Air Force Intelligence Analysis: Volume 1, Findings and Recommendations*, Santa Monica, Calif.: RAND Corporation, RR-A341-1, 2021.

² Lance Menthe, Amado Cordova, Carl Rhodes, Rachel Costello, and Jeffrey Sullivan, *The Future of Air Force Motion Imagery Exploitation: Lessons from the Commercial World*, Santa Monica, Calif.: RAND Corporation, TR-1133-AF, 2012; and Amado Cordova, Lindsay D. Millard, Lance Menthe, Robert A. Guffey, and Carl Rhodes, *Motion Imagery Processing and Exploitation (MIPE)*, Santa Monica, Calif.: RAND Corporation, RR-154-AF, 2013.

³ Lance Menthe, Amado Cordova, Elliot Axelband, Lindsay D. Millard, Abbie Tingstad, Endy M. Daehner, Kirsten M. Keller, and John Langley, *Technologies and Processes for Automating Processing, Exploitation, and Dissemination*, Santa Monica, Calif.: RAND Corporation, 2015b, Not available to the general public.

2015, PAF looked at AF DCGS as part of the larger intelligence community (IC), defined the spectrum of activities for synthesizing intelligence, and made broad recommendations for automation.⁴ In FY 2017, PAF looked at the process for assessing AF DCGS intelligence and made recommendations to improve data structures, many of which are echoed in the most recent research.⁵

As intelligence needs and expectations change, analytic requirements will change along with them. However, some common through lines continue to emerge. For example, in the final report for the FY 2012 project on automating Air Force PED, PAF researchers said:

When it comes to PED, the Air Force's most valuable asset is, and will remain, its force of trained human analysts. Instead of trying to replicate by machine what its analysts already do well, the Air Force should invest in technologies that help analysts do more. In the near term, the emphasis should be on automating specific tasks to meet immediate needs. In the middle term, the Air Force should acquire new PED capabilities. In the long term, the AF DCGS should transition to an alternative workflow and organizational construct better suited to take advantage of these capabilities.⁶

This insight is the basis for recommendations made throughout PAF's FY 2018 research on artificial intelligence (AI)/machine learning (ML), as detailed in Volume 1 of this series and this report. However, two factors have changed since the former and present research was conducted. First, technologies that can fully automate parts of the analytic process—specifically, the AI/ML methods described in Chapter 4 of this report—are at last starting to mature. Many of the capabilities that PAF previously recommended for acquisition only in the “middle term” can (and should) be sought now. Second, the demand for AF DCGS support has broadened and is expected to continue to expand beyond the laser-like focus on COIN/CT operations experienced after the September 11, 2001, terrorist attacks. The AF DCGS may need to support a wider variety of more-challenging intelligence problems in the higher-end threat environments envisioned by the 2018 National Defense Strategy. This development places a greater premium on AF DCGS operational agility going forward; consequently, the FY 2018 research includes agility as one of the enduring challenges to be addressed in the future.

⁴ Brien Alkire, Abbie Tingstad, Dale Benedetti, Amado Cordova, Irina Elena Danescu, William Fry, D. Scott George, Lawrence M. Hanser, Lance Menthe, Erik Nemeth, David Ochmanek, Julia Pollak, Jessie Riposo, Timothy William James Smith, and Alexander Stephenson, *Leveraging the Past to Prepare for the Future of Air Force Intelligence Analysis*, Santa Monica, Calif.: RAND Corporation, RR-1330-AF, 2016.

⁵ Abbie Tingstad, Dahlia Anne Goldfeld, Lance Menthe, Robert A. Guffey, Zachary Haldeman, Krista S. Langeland, Amado Cordova, Elizabeth M. Waina, and Balys Gintautas, *Assessing the Value of Intelligence Collected by U.S. Air Force Airborne Intelligence, Surveillance, and Reconnaissance Platforms*, Santa Monica, Calif.: RAND Corporation, RR-2742-AF, forthcoming.

⁶ Menthe et al., 2015b.

Scope

We recommend tools, technologies, and processes to address the growing demand for AF DCGS support, specifically for the work of the intelligence squadrons within the 480th Intelligence, Surveillance, and Reconnaissance Wing (ISRW) and associated units. The recommendations extend to core Distributed Ground Station (DGS) sites, Air National Guard (ANG) sites, and Distributed Mission Site (DMS) locations. Although other wings and organizations within the Air Force, such as the 70th ISRW, 55th ISRW, and National Air and Space Intelligence Center (NASIC), also perform time-sensitive intelligence analysis, we did not examine their processes. However, the results and methods from this project also should be of interest to them.

We consider analysis of the following types of collection: high-altitude imagery, including electro-optical (EO)/infrared (IR) images and synthetic aperture radar (SAR) images; motion imagery, including FMV and wide-area motion imagery (WAMI) collections; moving target indicator (MTI); electronic intelligence (ELINT); and SIGINT. Although the processes for analyzing less-common forms of imagery, such as coherent change detection⁷ or hyperspectral imagery (e.g., from the Airborne Cueing and Exploitation System–Hyperspectral [ACES-Hy] sensor⁸) often mirror standard processes for high-altitude imagery, we do not discuss them in detail.

In this research, we give special attention to short-term fixes to existing challenges and where AI/ML could be integrated into analytic processes in the coming years. We also consider how the AF DCGS could leverage the investments of partner organizations going forward. In addition to which tools and technologies can improve AF DCGS operations, we also look at AF DCGS management processes, how the AF DCGS onboards new tools and technologies, and how AF DCGS fosters innovation generally. We show that investments in AI/ML require supporting improvements in several areas. This research should be used to guide these investments.

In Volume 1 and this report, we focus on the evolution of the AF DCGS toward a fully integrated, multi-intelligence (INT) organization. Specific system recommendations for each INT are provided in a separate restricted volume.⁹ The restricted volume also contains a catalog of the many tools, networks, and systems in use at the AF DCGS.

⁷ See General Atomics Aeronautical, “Lynx Multi-Mode Radar,” website, undated.

⁸ Amy Butler, “Eyes Wide Open,” Aviation Week Network, September 19, 2011a; and Amy Butler, “USAF Turns to Hyperspectral Sensors in Afghanistan,” Aviation Week Network, September 19, 2011b.

⁹ Lance Menthe, Dahlia Anne Goldfeld, Sherrill Lingel, Abbie Tingstad, and Anne Stickells, *Technology Innovation and the Future of Air Force Intelligence Analysis: Volume 3, Technical Assessment of Data Flow Maps*, Santa Monica, Calif.: RAND Corporation, forthcoming. Not available to the general public.

Terminology

For consistency with earlier PAF work and other publications, we continue to use several legacy terms in this report. We discuss them here for clarity, along with some of the newer terms that are replacing them.

Processing, Exploitation, and Dissemination

Joint Publication 2-01 now defines the *intelligence process* as consisting of six operations: (1) planning and direction, (2) collection, (3) processing and exploitation, (4) analysis and production, (5) dissemination and integration, and (6) evaluation and feedback.¹⁰ The older abbreviation *PED* refers to step 3, the first phase (or time-dominant analysis) portion of step 4, and the dissemination aspect of step 5. However, *PED* remains in common currency in the AF DCGS in part because it describes the specific portions of the intelligence process that AF DCGS currently conducts. The term has been in common usage for long enough that it is also sometimes used as a verb, although we do not do so here.

Related abbreviations are *CPED*, where the “C” stands for collection, and *TCPED*, where the “T” stands for tasking, a subset of planning and direction. There is also the broader abbreviation *PCPAD*, which means planning and direction, collection, processing and exploitation, analysis and production, and dissemination.

Sense, Identify, Attribute, and Share

A new term—*sense, identify, attribute, and share* (SIAS)—was introduced in mid-2018 to reflect a nonlinear, information-age view of the analysis process, as opposed to the term *PED*, which is more rooted in an industrial production line view of the process.¹¹ Unlike traditional, stove-piped models of analysis that are designed to answer basic intelligence questions from airborne collectors, SIAS emphasizes fusing data from all sources to answer more-advanced intelligence questions and distribute that information widely. Although this project was completed before this new paradigm could be fully defined, the recommendations to automate basic tasks and rebalance AF DCGS efforts toward solving intelligence problems and supporting analysis are fully consistent with this direction.

Collection Disciplines

Following common practice, we refer to the various collection disciplines (or forms of intelligence gathering) as *INTs*. Over the past decade or so, the traditional term *imagery*

¹⁰ Joint Publication 2-01, *Joint and National Intelligence Support to Military Operations*, Washington, D.C., July 5, 2017.

¹¹ John A. Tirpak, “‘PED Is Dead’: ISR Roadmap Reaches Long for New Tech,” *Air Force Magazine*, August 2, 2018.

intelligence (IMINT) has been slowly phased out in the IC in favor of *geospatial intelligence* (GEOINT), which is defined as “the exploitation and analysis of imagery and geospatial information to describe, assess, and visually depict physical features and geographically referenced activities on the earth.”¹² GEOINT also includes imagery-derived measurement and signatures intelligence (MASINT) and related disciplines.¹³ To a lesser extent, the term *imagery-intelligence analyst* (IA) is being phased out in favor of *geospatial-intelligence analyst* (GA), and the term *imagery-intelligence reports editor* (IRE) is being phased out in favor of *geospatial-intelligence reports editor* (GRE).¹⁴ However, the term *imagery-intelligence mission supervisor* (IMS) remains standard.¹⁵ We use GEOINT, GA, and GRE in this series of reports.

Finally, a relatively new discipline for the AF DCGS is open-source intelligence (OSINT). We define *OSINT* as “publicly available information that has been discovered, determined to be of intelligence value, and disseminated by a member of the IC.”¹⁶

Phases of Analysis

Throughout this research, we refer to *phases of analysis*.¹⁷ Definitions of these phases vary throughout the IC. We outline what we mean by these phases in Table 1.1. For the most part, we define the phases in terms of timeline. These are our definitions, which reflect common use and practice in the AF DCGS community. It should be noted that the timelines under which these processes are completed can be much faster than shown during crisis situations.

Note that, in our research, we use the term *all-source analyst* as NASIC does to denote “national experts on threats that span air, space, and cyberspace domains.”¹⁸ As described in Chapter 2 of this report, an all-source analyst represents the deepest form of synthesis. (There might be some confusion here because the Air Force also names its most-generic intelligence career field, Air Force Specialty Code [AFSC] 1N0, an “all-source intelligence analyst.”¹⁹)

¹² Office of the National Geospatial-Intelligence Agency Historian, *The Advent of the National Geospatial-Intelligence Agency*, Washington, D.C., September 2011, p. 3. The National Imagery and Mapping Agency, founded in 1996, changed its name in 2003 to the National Geospatial-Intelligence Agency (NGA) to reflect this larger view.

¹³ National System for Geospatial-Intelligence, *Geospatial Intelligence (GEOINT) Basic Doctrine*, Publication 1.0, Springfield, Va., April 2018.

¹⁴ Usage is inconsistent. Different documents and sites prefer different terms.

¹⁵ IMS likely remains standard because the abbreviation for GMS, which stands for ground mission supervisor, is already taken.

¹⁶ Heather J. Williams and Ilana Blum, *Defining Second Generation Open Source Intelligence (OSINT) for the Defense Enterprise*, Santa Monica, Calif.: RAND Corporation, RR-1964-OSD, 2018, p. 8.

¹⁷ These may also be called phases of exploitation.

¹⁸ National Air and Space Intelligence Center, “About Us: National Air and Space Intelligence Center,” webpage, May 2018.

¹⁹ U.S. Air Force, *Air Force Specialty Code 1N0X1: All Source Intelligence Analyst Career Field Education and Training Plan*, Washington, D.C.: Department of the Air Force, September 26, 2016, p. 18.

Table 1.1. Phases of Analysis

Phase	Typical Activity	Timeline
0	Determine whether the collection is useful or must be retaken. Immediate threat warnings.	Seconds to minutes
1	Basic information of “who, what, when, and where.” Initial judgments of “why and how.”	Minutes to hours
1.5	Multisource correlation and association to provide context for operational needs. Revised judgments of “why and how.”	Hours to days
2	Multi-INT fusion to derive new information and answer intelligence questions. Forensic analysis of older data.	Days to months
3+	All-source analysis.	Months to years

We also sometimes refer to *time-dominant* versus *content-driven* analysis. *Time-dominant analysis* is defined as “tradecraft focused on rapid discovery by correlating what is new with what is known.” *Content-driven analysis* is “expository” and “places less emphasis on rapidity or the data source and greater emphasis on . . . analytic depth and explanatory narratives.”²⁰ The AF DCGS today formally describes its analytic objective as providing time-dominant analysis.²¹

Research Methodology

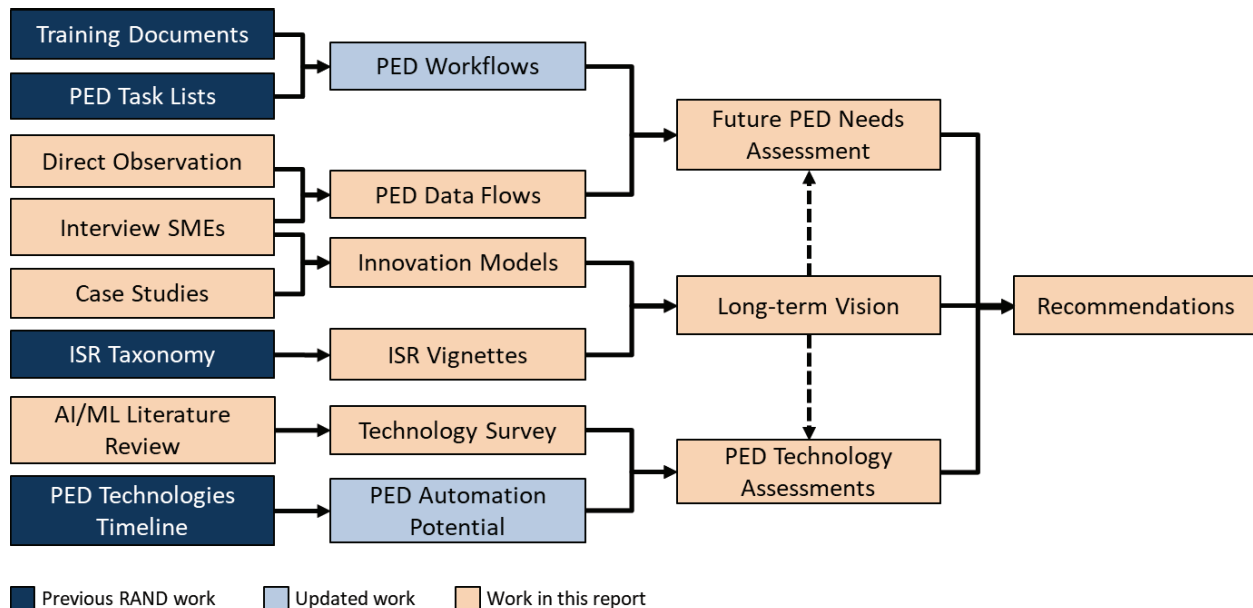
In our research, we employed multiple qualitative analysis methods, which we illustrate in Figure 1.1. For some portions of the analysis, we leveraged previous PAF work on PED tasks, training documents, and workflows, as well as a recently developed taxonomy for ISR missions.²² We reviewed training documents, including job-qualification standards and training task lists (TTLs) for all positions. New or updated analyses largely focused on understanding historical and current practices and challenges, including the case studies described in Chapter 3; assessing the state of the art in AI/ML, as described in Chapter 4; and identifying lessons for developing and fielding new technologies.

²⁰ The terms *time-dominant* and *content-driven* were first defined by Jason M. Brown and David Vernal, “Time-Dominant Fusion in a Complex World: Defining Time-Dominant Fusion and Its Interdependent Relationship with Airborne ISR Capabilities and Air Force DCGS,” *Trajectory*, November 11, 2014.

²¹ Air Combat Command Manual 14-401, *Air Force Distributed Common Ground System (DCGS) Training, Certification, and Quality Management*, Joint Base Langley-Eustis, Va., April 6, 2020.

²² Tingstad et al., forthcoming.

Figure 1.1. Project Approach



NOTE: SMEs = subject-matter experts.

The largest single effort in the project was the direct observation of PED processes, along with informal interviews with analysts.²³ In addition to multiple discussions with staff within Headquarters Air Force/A2 and Headquarters 480th ISRW, we conducted many site visits and interviews to update our knowledge of PED challenges: specifically, with DGS-1, DGS-2, DGS-3, DGS-5, DGS-IN, DMS-GA, and DMS-HI.²⁴ We also visited the 11th Special Operations Intelligence Squadron and spoke with SMEs at NASIC/Global Exploitation Intelligence Group (GX), the National Security Agency (NSA), and the National Reconnaissance Office (NRO). Finally, we held informal discussions on nonmilitary AI/ML issues with SMEs from Netflix and Google.

In our visits to AF DCGS sites, we updated our prior understanding of analyst workflows derived from the TTLs for each AF DCGS crew position. TTLs outline what skills and knowledge an AF DCGS crew member must demonstrate to be qualified to perform a specific position. For example, crew members must generally understand the types of Air Force platforms that collect intelligence, know how to use their workstations and software, be familiar with common dissemination pathways, and be able to explain why different aspects of a pre-mission brief (PMB) are important. However, initial interviews with analysts, as well as previous PAF work, revealed that TTLs neither include all the steps that analysts take in conducting their

²³ We did not visit organizations to discuss all-source analysis because this was outside the scope of the AF DCGS.

²⁴ The abbreviations for states are IN = Indiana, GA = Georgia, and HI = Hawaii. We did not visit the final core site, DGS-4, because of resource constraints; instead, we drew from previous experience conducting multiday interviews at this site. See Alkire et al., 2016.

work nor adequately capture PED crew teamwork.²⁵ Because of this, we developed our own descriptions of the analysis roles within the AF DCGS, as discussed in Chapter 2 of this report. We also captured our knowledge of current PED processes in what we call *data flow maps*, which we describe in the next section.

Using our assessment of future needs and our understanding of current and potential future AI/ML capabilities, we produced a long-term vision for improving AF DCGS operations. The vision includes both short-term recommendations that can be implemented using today's capabilities and longer-term recommendations that take advantage of evolving capabilities as they mature. A major conclusion of this research is that although technology innovation has great potential, the benefits cannot be realized without significant attention to the human factors—the organization, training, and even culture that determine whether new technologies are accepted and effective. Consequently, we convened a group of RAND experts to review our preliminary recommendations on matters of personnel, training, and organization. In this group elicitation, we presented the collected observations from our site visits and several preliminary recommendations. The panel members provided suggestions regarding implementation and other considerations that were included in our final recommendations.

About Data Flow Maps

As just mentioned, we examined TTLs for each collection discipline, as well as for multi-INT processes, and conducted on-site observation and interviews to understand the workflow of the PED crews. There are many ways to view this, but, for our purposes, we focused on how analysts manipulate data to create the final products that are then disseminated in different ways. Using this analysis, we created data flow maps for each INT.

Data flow maps are flowcharts or “circuit board” illustrations of how data are transformed through analysis. They show the sequence in which software tools are used, how data are passed from one computer system to another, and where key analytic tasks must occur. In this project, we use data flow maps to show how PED crews transform sensor data into intelligence products and services and disseminate them to the warfighter. This approach is designed to illuminate dependencies and identify how new tools would affect the overall analysis process. By explicitly showing these interconnections, we can also see how alleviating a bottleneck in one part of a PED crew's work may lead to a secondary bottleneck elsewhere. It should be noted, however, that these maps do not display the full communications network architecture and are not intended to show every task performed by the PED crew.

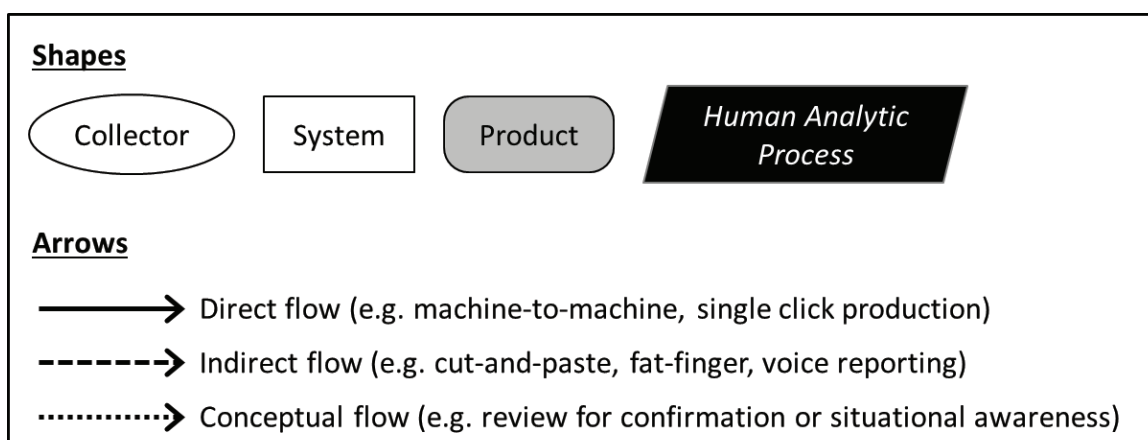
The maps were constructed via extensive observation and interviews with analysts at several AF DCGS sites. We then reviewed each initial map with at least three different experienced analysts to make final adjustments. In these maps, we do not model PED crew positions

²⁵ See Alkire et al., 2016; Menthe et al., 2012; and Menthe et al., 2015b.

explicitly but instead focus on the work of the team. The result is a representation that is independent of crew position and should remain relevant as the AF DCGS changes those positions and moves toward the SIAS paradigm.

Figure 1.2 shows the data flow maps that distinguish between four types of objects (represented as the basic shapes) or places where data can dwell: *collectors*, which gather the data in the first place; *systems*, a broad category that includes hardware, software, and databases; *products*, which include both final products and intermediate formats, such as lines of text, .kml files, or .pptx files; and *human analytic processes*, which are steps in the process where analysts manipulate data and gain understanding.

Figure 1.2. Data Flow Map Symbols—Current Processes



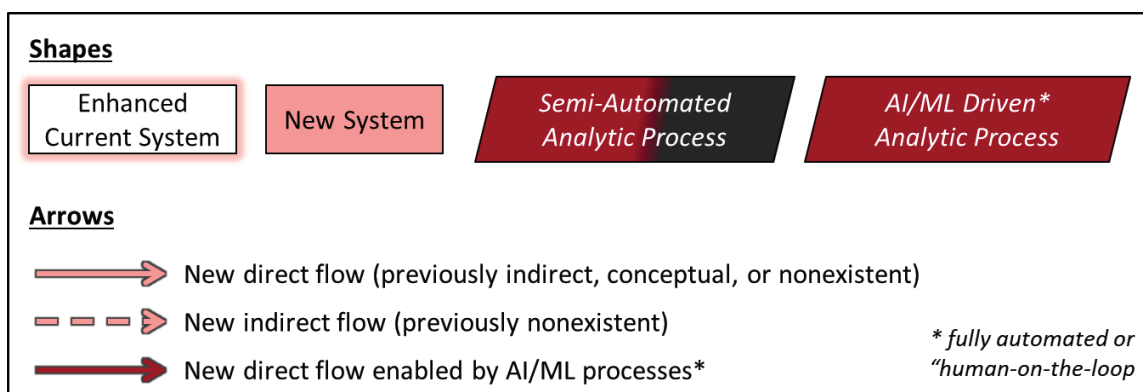
There are also three different types of connectors indicating how data flow between objects: *direct*, *indirect*, and *conceptual*. Each movement or manipulation of the data is represented using one of these connectors. *Direct flow* refers to either machine-to-machine flow or movement that requires very little human intervention, such as pushing a single button. *Indirect flow* requires significant human intervention, such as filling out a complicated online form. We focused on this distinction because, using previous PAF research,²⁶ we anticipated that much of the “low-hanging fruit” for automation would involve converting indirect flow to direct flow, reducing the burden on analysts to manually enter data from one system to another. Finally, *conceptual flow* is where analysts review extra data to maintain situational awareness or to confirm their exploitation, but they do not generally copy that information forward. The bulk of an analyst’s time may involve such an activity, which can provide crucial context but does not usually leave a clear trail.

²⁶ Tingstad, forthcoming; Menthe et al., 2015b.

In Chapter 2 of this report, we present a generic data flow map that is broadly representative of current practice. In the restricted volume,²⁷ we provide more-specific data flow maps for high-altitude imagery, including EO/IR and SAR images; motion imagery, including FMV and WAMI collections; MTI; ELINT; and SIGINT.

We further use data flow maps to illustrate some of our recommended improvements to AF DCGS tools and processes (see Chapters 5 and 6 of this report). Figure 1.3 shows the types of changes identified on those maps.

Figure 1.3. Data Flow Map Symbols—Potential Future Entities and Connections



For some recommendations, a system is merely enhanced, such as by adding a field to a database. In other cases, a new system is required, which may be a new software program or just a script that performs a specific function. We also introduce semiautomated and fully automated analytic processes. Finally, in addition to pink arrows indicating new direct and indirect flows, a dark red arrow indicates a flow that is governed by AI/ML processes (i.e., places where a human may be *on* the loop but where human intervention is no longer required).²⁸

In the restricted volume, we use color coding to indicate alternative pathways where different DGS sites use different systems or where there is a commonly used alternative procedure.²⁹ In part because they support different combatant commands (CCMDs) and must share data with different partners, the analytic processes at different sites within the AF DCGS are similar but not always identical. There are some additional symbols representing different types of systems. Where we introduce changes, we sometimes also present more than one alternative, depending on which proves more technically feasible.

²⁷ Menthe et al., forthcoming.

²⁸ *Human in-the-loop* refers to an otherwise automated process with a step that must be performed by a human. The process effectively pauses until the human acts. *Human on-the-loop* is an automated process that a human may choose to pause or override, but where the loop will otherwise proceed without human intervention.

²⁹ Menthe et al., forthcoming.

In addition to illustrating PAF's recommendations, the data flow maps in this volume and the restricted volume should be a useful reference for the AF DCGS to understand its own roles, processes, organization, challenges, and solutions.

2. Overview of the AF DCGS Today

In Volume 1, we briefly describe how the AF DCGS operates, with particular emphasis on present challenges.³⁰ This chapter provides a more in-depth discussion of how AF DCGS evolved and functions today. We first describe how the AF DCGS evolved into its present form amid the massive increase in data supply and warfighter demand. We then discuss how the AF DCGS is organized, both globally and within a particular site. We present a taxonomy of roles that describes various crew positions and show how those roles interact within the larger data flow map. This description, and especially the data flow map, form the baseline for assessing challenges and potential solutions.

Evolution of the AF DCGS and the Growth in Supply and Demand

History helps put the current growth in context. The AF DCGS was established in 1996.³¹ It was the direct successor to the Contingency Airborne Reconnaissance System, established in 1992 at what was then Langley Air Force Base (AFB). Originally created to conduct PED for U-2 overflight missions, the AF DCGS expanded greatly when it began to provide PED for FMV and other data collected from remotely piloted aircraft (RPA) flying in support of Operation Enduring Freedom (OEF) and Operation Iraqi Freedom (OIF)—and has continued to do so for subsequent operations.³² The AF DCGS now conducts PED around the clock for more than 50 ISR sorties every day.³³

The Air Force ISR community underwent an intense evolution after 9/11. OEF and OIF drove the AF DCGS to retrain and reorganize around primarily COIN/CT missions. Much was gained as the AF DCGS grew. Several characteristics of COIN/CT operations affected the development of new systems and the tactics, techniques, and procedures (TTPs) for the AF DCGS: the rapid acquisition of RPA to provide 24/7 near-real-time intelligence in permissive

³⁰ Menthe et al., 2021.

³¹ The AF DCGS is formally designated AN/GSQ-272 SENTINEL. The AF DCGS was initially envisioned as part of a family of U.S. Department of Defense (DoD) systems, including the U.S. Army's DCGS-A, U.S. Navy's DCGS-N, and U.S. Marine Corps' DCGS-MC. In practice, however, the AF DCGS has become a *much* larger, effectively separate system. A note on pronunciation: The Air Force's DCGS is spelled out ("dee cee gee ess"), whereas the others are pronounced "dee-sigs."

³² Although OEF and OIF greatly expanded their use, the first deployments of what was then the RQ-1A Predator were in support of NATO air operations in the Balkans in 1995–1996 and 1999. See Defense Airborne Reconnaissance Office, *UAV Annual Report FY 1996*, Washington, D.C.: Office of the Under Secretary of Defense, Acquisition and Technology, November 6, 1996; and Benjamin S. Lambeth, *NATO's Air War for Kosovo: A Strategic and Operational Assessment*, Santa Monica, Calif.: RAND Corporation, MR-1365-AF, 2001.

³³ U.S. Air Force, "Air Force Distributed Common Ground System," webpage, October 13, 2015b.

environments; the development of new capabilities that enable reach-back;³⁴ and the focus on prosecuting fleeting targets and target networks, especially for CT missions.³⁵

The Air Force took the initiative to support these operations, including engaging in near-continuous surge operations for many years. This required the onboarding of new personnel, systems, tools, and processes—essentially “building the aircraft in flight”—while shouldering responsibility for lives on the ground. The result has been an Air Force PED enterprise that has expanded quickly by almost any measure. Nevertheless, wartime conditions do not always allow for careful long-term planning, and thus PED technologies have been developed and interconnected in novel, sometimes ad hoc ways.

Complicating this burden is that expectations have grown as well. The AF DCGS is, in some sense, a victim of its own success. Warfighters in COIN/CT operations have grown accustomed to tailored PED support, which, in turn, has raised expectations about PED responsiveness and availability in general. Furthermore, the targets themselves have become more challenging because the complexity and mobility of targets—both of which can make PED more difficult—have been on the rise among U.S. adversaries and other potential challengers.

Moreover, with new ISR capabilities come new demands. For example, when the resolution of FMV cameras in the original MQ-1 Predators was upgraded, some believed that the number of analysts required would go down because it would be easier to fulfill requests for essential elements of information (EEIs). But of course, that is not what happened; the questions only got harder. Instead of just detecting vehicles, analysts were now asked to provide details about color and model. Analysts continue to glean as much information as possible from each pixel, as they have done for still imagery; FMV only creates more pixels to be analyzed. As one expert noted:

An anecdote familiar to many senior leaders concerns a numbered air force commander’s use of a single slide in 2007 to accentuate a point about ISR. This slide (used effectively in many meetings) depicted a startling contrast between the growth in ISR [Combat Air Patrols (CAPs)] and a rough order-of-magnitude measure of combatant command and national ISR requirements. Specifically, for every increase in ISR capability (CAPs increase), the documented needs grew at a greater, expanding rate. This fact underscored what we previously treated as a useful exaggeration: the never-ending appetite for ISR.³⁶

The sheer amount of data collected remains a primary challenge for the AF DCGS today. The number of ISR missions that the AF DCGS supports skyrocketed with OEF and OIF, increasing

³⁴ *Reach-back* is “the process of obtaining products, services, and applications, or forces, or equipment, or material from organizations that are not forward deployed” (U.S. Joint Chiefs of Staff, Joint Doctrine Division, DOD Dictionary of Military and Associated Terms, Washington, D.C., June 2018).

³⁵ Some capabilities were also lost. Gen Herbert J. Carlisle noted that insufficient emphasis was placed on electronic warfare and the cyber domain for a period of time because “[in] the air domain, it was less of a factor in Afghanistan and Iraq. We didn’t need to” (Mark Pomerleau, “Carlisle: Overworked Airmen Can’t Train for Future Threats,” *Defense Systems*, September 18, 2015).

³⁶ Jon Kimminau, “A Culminating Point for Air Force Intelligence, Surveillance, and Reconnaissance,” *Air and Space Power Journal*, Vol. 26, No. 6, November–December 2012, pp. 119–120.

almost twentyfold in 14 years.³⁷ As early as 2009, then–Air Force Deputy Chief of Staff for ISR Lt Gen David A. Deptula warned that, “in the not too distant future,” the Air Force would be “swimming in sensors and drowning in data.”³⁸ A PAF report from an FY 2010 project cautioned that, if the Air Force did not change how it conducted PED, the AF DCGS would need more than 100,000 analysts by 2020.³⁹

The direst predictions of data overload have yet to materialize, largely because the WAMI sensors that were under development at the time, notably Gorgon Stare,⁴⁰ have not been adopted as widely as expected. The Air Force has so far acquired only eight Gorgon Stare sensors for its fleet of MQ-9 Reapers. This is in part because, when so equipped, Reapers cannot also carry weapons, but also because Gorgon Stare resolution and frame rates are substantially lower than for standard FMV, which limits operational utility. Current plans no longer call for widespread deployment of these sensors across today’s RPA fleet.⁴¹ Furthermore, the Air Force chose to sidestep much of the PED problem posed by WAMI sensors by leaving the vast majority of the data they collect unexamined.⁴² Finally, fears that the AF DCGS would need to support 90 or more continuous ISR sorties per day subsided when OEF and OIF wound down in the mid-2010s.⁴³

Nevertheless, the amount of data collected by Air Force platforms continues to grow as sensors improve in resolution and sensitivity—and this is not just a feature of motion imagery but also for collections of all types, including SIGINT, SAR imagery, and MTI. In 2011, the Air Force estimated that the AF DCGS ingested 700 gigabytes per day;⁴⁴ the most recently available

³⁷ Timothy D. Haugh and Douglas W. Leonard, “Improving Outcomes: Intelligence, Surveillance, and Reconnaissance Assessment,” *Air and Space Power Journal*, Vol. 31, No. 4, Winter 2017.

³⁸ David A. Deptula, keynote speech, C4ISR Journal Conference, Arlington, Va., October 2009; and David A. Deptula, “Air Force ISR in a Changing World: Changing Paradigms While Optimizing ‘Low Density’ to Meet ‘High Demand,’” in Keith Brent, ed., *The Art of Air Power: Proceedings of the Royal Australian Air Force Air Power Conference*, Canberra, Australia: Commonwealth of Australia, March 30, 2010.

³⁹ Menthe et al., 2012.

⁴⁰ Michael Hoffman, “Gorgon’s Gaze Set for Fall in Afghanistan,” *Air Force Times*, June 13, 2010.

⁴¹ Per Rachel Cohen, “Gorgon Stare to Receive BLOS Upgrades While Air Force Explores Replacement,” *Inside Defense*, April 6, 2018:

Gorgon Stare, the Air Force’s sensor program of record for wide-area motion imagery mounted on MQ-9 remotely piloted aircraft, will receive a limited slate of upgrades in the near future but isn’t currently expected to grow across the Reaper fleet, the service recently reported to Congress.

⁴² Some say that only 15 percent of the footage has been examined; we believe this to be a generous estimate (Yasmin Tadjdeh, “Algorithmic Warfare: Google Versus the Pentagon, the Fallout,” *National Defense Magazine*, August 2, 2018).

⁴³ William Giannetti, “A Commonsense Approach to Intelligence, Surveillance, and Reconnaissance Operations,” *Air and Space Power Journal*, Vol. 30, No. 3, Fall 2016.

⁴⁴ Unclassified estimate from an earlier version of an AF DCGS fact sheet from 2011 (cited in Menthe et al., 2015b).

estimate is 20 terabytes per day, a nearly thirtyfold increase.⁴⁵ Even without WAMI, the rate at which data flow through the AF DCGS continues to grow rapidly.

If the Air Force chooses to invest in more-capable WAMI sensors, the data rate could jump further. “Increment 2” of the Gorgon Stare program,⁴⁶ using the Defense Advanced Research Project Agency’s (DARPA’s) Autonomous Real-Time Ground Ubiquitous Surveillance Imaging System (ARGUS-IS), can collect 10 petabytes of raw data per day—which, even with today’s thousandfold video-compression techniques, would be staggering.⁴⁷ Moreover, the recent explosion in OSINT and PAI has the potential to add vast amounts of new data, depending on the extent to which they are incorporated into AF DCGS processes. The Air Force must anticipate that data-collection rates will continue to grow for the foreseeable future, and it should account for the risk that a new technology could generate a sudden, discontinuous jump in analytic demand.

Global Organization

Meeting the ever-expanding supply of data and demand for ISR PED requires a vast global capability. The 480th ISRW is the lead wing responsible for executing AF DCGS operations. All active-duty units involved in the AF DCGS, including the 480th ISRW, are part of the 25th Air Force, under the Air Combat Command (ACC).⁴⁸ The core AF DCGS locations are the five numbered DGS sites: DGS-1, at Joint Base Eustis-Langley, led by the 497th ISRG (ISRG); DGS-2, at Beale AFB, led by the 548th ISRG; DGS-3, at Osan Air Base, led by the 694th ISRG; DGS-4, at Ramstein Air Base, led by the 693rd ISRG; and DGS-5, at Joint Base Pearl Harbor–Hickham, led by the 692nd ISRG.⁴⁹ These sites are regionally aligned and supported by several

⁴⁵ U.S. Air Force, 2015b.

⁴⁶ Loren Thompson, “Air Force’s Secret ‘Gorgon Stare’ Program Leaves Terrorists Nowhere to Hide,” *Forbes*, April 10, 2015.

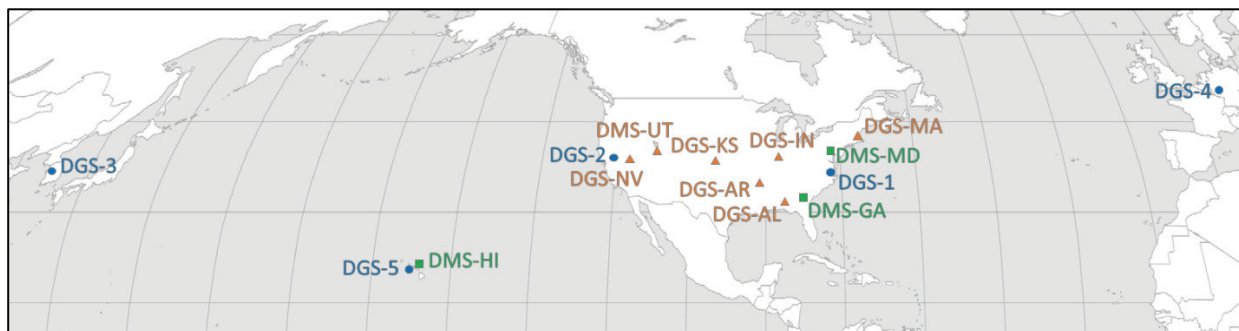
⁴⁷ See Brian Dodson, “DARPA’s New 1.8-Megapixel Camera is a Super High-Resolution Eye in the Sky,” *PBS*, February 11, 2013; and David Hambling, “New Army Camera Promises Super-Wide Surveillance,” *Wired*, August 19, 2009.

⁴⁸ Before the 25th Air Force was stood up in 2014, the AF DCGS had been under the purview of the Air Force Intelligence, Surveillance, and Reconnaissance Agency, which reported to Air Force/A2. (The agency essentially became the 25th Air Force under ACC.) See Wayne Amann, “Former AF ISR Agency Now Numbered Air Force,” U.S. Air Force webpage, 25th Air Force Public Affairs, October 2, 2014.

⁴⁹ Sixteenth Air Force (Air Forces Cyber), “480th ISR Wing,” webpage, September 8, 2019.

ANG sites, designated by state.⁵⁰ Additional DMSs provide specialized capabilities.⁵¹ Figure 2.1 shows the major AF DCGS locations.

Figure 2.1. Major AF DCGS Sites



SOURCE: U.S. Air Force, “U-2S/TU-2S,” webpage, September 23, 2015a; U.S. Air Force, 2016.

Distributed Ground Stations Site

Within each DGS is a complex organization with different functions and roles. In our research, we focus on the two main parts of AF DCGS operations, conducted chiefly by intelligence squadrons: the operations floor (“ops floor”) and DCGS Analysis and Reporting Teams (DARTs).⁵²

Ops Floor

The ops floor at a DGS is primarily organized around skilled PED crews.⁵³ Crews include members serving in various roles, such as exploiting, analyzing, reporting, and supervising. The

⁵⁰ The main ANG sites are DGS-AL, DGS-AR, DGS-IN, DGS-KS, DGS-MA, DGS-NV, and DMS-UT. Others are colocated with active-duty AF DCGS locations. Although no longer formally part of the AF DCGS, the 11th Special Operations Intelligence Squadron also conducts PED for special operations forces. See Hurlburt Field, “11th Special Operations Intelligence Squadron,” webpage, March 28, 2017. Also, what was formerly the DMS-NASIC continues to provide PED-like support. See Air Force Intelligence, Surveillance, and Reconnaissance Agency Instruction 14-153, *Air Force Distributed Common Ground System (AF DCGS) Operations Procedures*, March 15, 2013, 480th ISR Wing Supplement, February 5, 2014.

⁵¹ DMS-GA, DMS-MD, and DMS-HI.

⁵² The AF DCGS comprises more than just intelligence squadrons. There are support squadrons that provide network communications and other logistics, support training and qualifications, and support collaboration with the IC (see Alexandre Montes, “AF NTI Training Streamlines Intel Airmen to Mission,” press release, 70th ISRW Public Affairs, Fort George G. Meade, Md., November 17, 2016). There are also innovation cells, technical directors, contract support, and leadership. We were fortunate to have the opportunity to speak to SMEs in all of them. We also note several areas in the final recommendations where assistance in some of these areas, including mission management and training, are needed.

⁵³ This is true for both active-duty and ANG DGS sites, although organization at DMS locations may differ. When we say “DGS sites,” this is generally shorthand for most AF DCGS locations.

term *crew* arose out of the sense that PED analysts should not be considered an added luxury for ISR missions but rather an integral part of an aircraft's ground crew or logistics tail. This is a particularly important consideration for *RPA*, a term that the Air Force uses instead of *unmanned aerial vehicle* or *unmanned aircraft system* to emphasize that these platforms require an extensive human crew offboard.⁵⁴ PED crews conduct Phase 0 and Phase 1 analysis (as defined in Chapter 1), meaning that they are the first to see the collections, deriving what information they can from them in a matter of seconds, minutes, or hours.

Since the inception of the AF DCGS, the organization of PED crews has been such that different crews on the ops floor have conducted PED separately for the data collected by individual sensors on ISR platforms. To a large extent, this is still true today, with the notable exception that all high-altitude imagery collected by the aircraft supported by a given DGS site may be exploited by a combined team on the ops floor (sometimes called a “super-crew”). Because PED crews are associated with both a collection discipline and a platform, they may be referred to in either way—for example, as an “FMV crew” or a “Reaper crew” (because personnel are generally qualified by INT, the INT name is more commonly used). When a platform carries more than one type of sensor, it is generally supported by more than one PED crew.⁵⁵ This is referred to as the *platform-centric* or *INT-centric* approach to organizing PED.⁵⁶ Depending on the size of the DGS site and the level of activity, there could be as many as a dozen PED crews on the ops floor or as few as one. Figure 2.2 shows a snapshot of part of an ops floor.

⁵⁴ At one time, for example, an estimated 192 people were required to support a 24-hour MQ-1 Predator combat air patrol, including mission control, launch and recovery, and PED elements. This large figure includes manpower factors to account for multiple shifts and absences. See Menthe et al., 2012.

⁵⁵ For example, the U-2S carries both imagery and SIGINT payloads. U.S. Air Force, 2015a.

⁵⁶ Menthe et al., 2012.

Figure 2.2. AF DCGS Operations Floor



SOURCE: U.S. Air Force, 2015b.

DCGS Analysis and Reporting Team

DARTs were created in 2007 to fuse data from multiple sensors.⁵⁷ Generally speaking, the DART conducts Phase 1.5 analysis (as defined in Chapter 1), meaning it creates fused, actionable intelligence within a matter of hours or days to provide context to meet operational needs. Prior to the DART, there was little ability for AF DCGS analysts to attain a broader view of the missions that they supported or to use the information they discovered in other ways. As one commander wrote:

In the past, Air Force DCGS analysts processing, exploiting, and disseminating intelligence from airborne collection had limited access to any intelligence other than that derived directly from one or two airborne platforms and associated sensors.⁵⁸

It took several years for the DART to come into its own, however. One reason for this is that the INT-centric division of labor between PED crews naturally favors a stovepiped intelligence process. This is not unique to the AF DCGS, of course; rather, it is a general feature of intelligence analysis: “Collection stovepipes emerge because the separate collection disciplines are often based on different technologies, are managed independently, and often are rivals to one

⁵⁷ Brown and Vernal, 2014.

⁵⁸ Brown and Vernal, 2014.

another.”⁵⁹ There are also advantages to INT-centric organization from an efficiency standpoint: It works well for training and scheduling personnel who share common skills.⁶⁰

The portfolio of the DART has changed considerably over the years. When certain DGS sites were in the throes of punishing surge conditions, the DART was reduced to a skeleton crew performing more-limited tasks, such as preparing PMBs because the analysts were needed to perform first-phase analysis on the ops floor.⁶¹ Now, however, the DART has expanded and its missions are formally recognized.

The flexibility in the DART’s structure is part of what enables it to be innovative, not just in terms of new tools but also in developing new TTPs. Some sites, for example, are experimenting with a mission analysis cell in the DART that is dedicated to conducting assessments of the PED process itself. In general, the DART investigates a variety of pressing intelligence questions that can be probed—wholly or in part—through Air Force airborne ISR platforms, in addition to supporting immediate operational needs related to mission support and threat warning. It also gives a glimpse of how the AF DCGS may evolve in coming years, as discussed in later chapters.

Crew Positions and Roles

Prior to the start of their mission, most PED crew members conduct common tasks associated with attending the PMB, such as fulfilling the pre-mission readiness checklist, initializing computer systems and programs, and participating in the debriefing that accompanies a crew changeover. During the mission, each crew member executes individual tasks—not necessarily sequentially—based on their crew position to accomplish their work, whether it be to support an ISR platform or to probe a relevant intelligence question. Although each crew position involves distinct tasks—tasks that can vary depending on operational circumstances—we find that most of them can be grouped usefully into four analysis roles: *exploiter*, *investigator*, *reporter*, and *supervisor*. Depending on the collection discipline, some crew members may perform more than one role: for example, an exploiter may also conduct analysis.⁶² Note also that different roles sometimes include similar tasks. These roles are summarized in Table 2.1 and are defined after the table.

⁵⁹ Mark M. Lowenthal, *Intelligence: From Secrets to Policy*, 5th ed., London: SAGE Press, 2012, p. 222.

⁶⁰ Menthe et al., 2012.

⁶¹ There was also concern at some point that the DART would conduct full Phase 2 intelligence and “compete” with NASIC, but these concerns no longer seem to arise. These days, few seem concerned that the appetite for intelligence can be so easily sated.

⁶² These categories of PED crew positions were developed by PAF by studying TTLs and through direct observation during site visits. They are not doctrinal.

Table 2.1. Analysis Roles

Role	Description	Sample Positions
Exploiter	Performs first phase analysis of ingested data. Answers fundamental questions about specific collections	GA, signals analyst, cryptologic operator (CO)
Investigator	Pursues intelligence questions. Fuses information. Conducts deeper levels of synthesis	Correlation analysis (CAN), mission support analyst (MSA), many DART members
Reporter	Prepares and stores intelligence products and warning messages	GRE, technical reporter
Supervisor	Oversees missions and personnel, clarifies tasks, performs quality control, conducts assessments	IMS, electronic intelligence mission supervisor (EMS), GMS

An *exploiter* is a crew member who focuses on answering basic “who, what, when, and where” questions through observation and standard data-manipulation techniques, such as that required for accurate geolocation of a signal. They may also address “why and how” questions to the extent that the requested EEs implicitly (or explicitly) require a determination of activity and intent. (This is part of Phase 0 and Phase 1 analysis.) A crew member in this role also often conducts background research to understand the context of the mission or improve their general situational awareness. Before interpreting the substance of the collection, an exploiter may use this auxiliary information to address broader questions, such as “Does the collection cover the correct area?” “Is it capable of providing the requested EEs?” “What else of significance is present?” “Is there anything here of immediate concern with respect to national security or protection of life?” Depending on the INT, they might dynamically retask or retune the collector using these preliminary assessments. Next, the exploiter will attempt to extract the requested EEs from the data at hand. They will then prepare the information in the appropriate product format for transmittal via active (e.g., email) or passive (e.g., databases) means. Note that, in many cases, the exploiter does not generate the final products—or at least, not all of them—but sends the information to a reporter to complete the process and disseminate the information. Examples of crew positions that fulfill the exploiter role are the GA⁶³ and the CO.

An *investigator* has a job description that is somewhat distinct from that of the exploiter. This type of analyst does not necessarily focus on active missions, although these might be taken into account; rather, they tackle broader problems. These problems can take the form of key intelligence questions that may require days (or months) of collection to answer, or they can focus on improving the quality of the missions themselves. In performing the former type of analysis, the investigator will typically draw from multiple sources to understand long-term patterns, develop context, and interpret activity over a long period. As opposed to the exploiter, who might state that “there is a red car,” an investigator could say “the preponderance of

⁶³ Also known as an *imagery analyst*.

evidence suggests that the red car belongs to a terrorist leader.” Another type of investigator could be employed to look critically at mission data to provide feedback on how the collection or exploitation might be better conducted in the future. This feedback might affect human as well as automated decisionmaking. Crew positions in the investigator role may also prepare reports that might be used to provide context for others, inform decisionmakers, or be stored to help build the case toward answering key intelligence questions down the road. Finally, some investigators work to enhance the reporting from active missions by adding context. Some examples of crew positions that fulfill the investigator role are the CAN and the MSA. In some cases, an MSA or a CAN might support more than one PED crew.

A *reporter* prepares intelligence products. Some reporters create finished products at the end of a shift, while others focus on providing threat warnings. Reporters must be familiar with how information needs to be transmitted, to whom, and via what means. They must also be familiar with the various messaging formats and systems necessary to transmit the information. Some examples of crew positions that fulfill the reporter role are the GRE, product reporter, technical reporter, and threat analyst. Many of these positions also have additional duties.

Supervisors can be present at many levels at a DGS site. Each PED crew has a supervisor—who may supervise more than one crew at a time—and there is also a mission operations commander who oversees the entire ops floor or some part of the ops floor who is dedicated to a subset of missions related to a common area, operation, or theme. Supervisors provide direction, such as determining which requirements each analyst will fulfill, assigning daily tasks, and enforcing quality control (e.g., ensuring that a report meets supported unit and legal standards). Sometimes, the quality-control process halts the workflow because the analyst is waiting for their assessment to be checked. Supervisors may also be responsible for communicating with the original requester of an EEI to clarify its intent. Some examples of crew positions that fulfill the supervisor role are the IMS, GMS, the EMS, and mission operations commander.

Data Flow and AF DCGS Processes

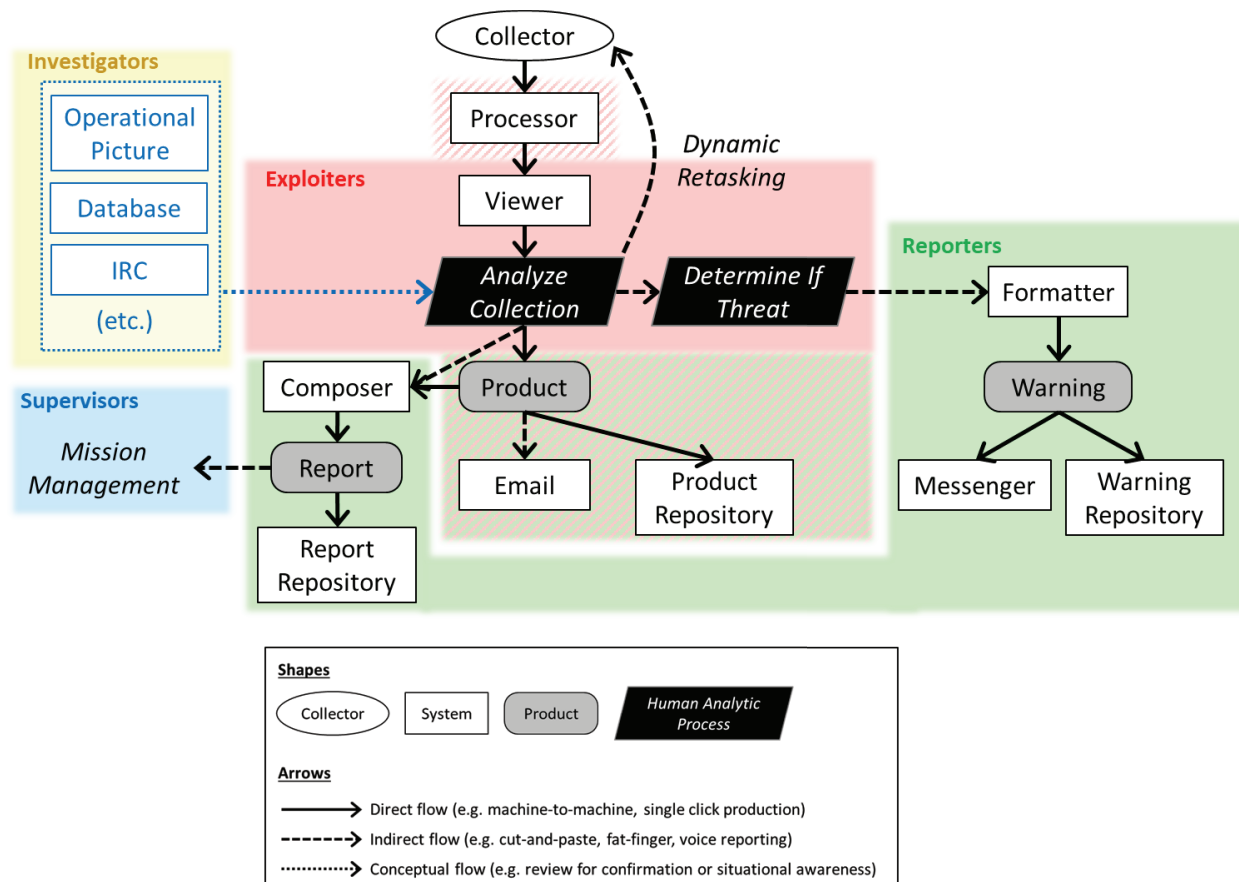
As discussed in Chapter 1 of this report, we developed data flow maps to depict current and potential future AF DCGS processes and how the various roles and supporting data infrastructure interact with one another. We discuss data flows for the ops floor and the DART in the next sections.

Ops Floor

Data flow maps are different for each INT. However, it is possible to construct a generic data flow map to illustrate the key points. Figure 2.3 shows a generic data flow for the ops floor, with colored underlays to indicate the four analysis roles discussed in the previous section. Note that the divisions between these roles are approximate and sometimes overlap. Additionally, the

exploiter often performs some of the investigator role, particularly while waiting for the platform that they are supporting to arrive on station.

Figure 2.3. Generic Data Flow Map—Current



NOTE: IRC = internet relay chat.

The map is designed to be read from top to bottom: the arrows are arranged to show the process beginning with collection,⁶⁴ moving through processing and exploitation, and ending with dissemination and storage.

After the data are collected, they are processed to remove noise and transform them into a more standard format that is suitable for transmission and storage. Depending on the sensor, this might occur onboard the aircraft or at the AF DCGS itself, which is why we mark the process as only partly in the exploiter role. But where it requires a human in-the-loop or on-the-loop, the exploiter is responsible for this task. Where the processing must be assisted by the human, the processor and viewer are usually two aspects of the same tool.

⁶⁴ Collection is not, of course, part of the PED process, but we include it to provide a familiar anchor to the diagram.

Once the data are in a human-readable format within the AF DCGS, the exploiters can analyze the collection. This includes determining whether the collection is useless (e.g., an image that shows only clouds) or potentially threatening (e.g., an unexpected emission on a reserved frequency). In those cases, the exploiter can initiate dynamic retasking or cross-cueing of sensors, either by retuning the sensor directly or communicating with the pilot and sensor operator to request the changes (e.g., adding a new target to the collection deck). If the exploiter determines that what they have found may constitute an immediate threat to the aircraft in the vicinity, they initiate the threat-warning process: They send the information to a reporter, who issues alerts in the proper formats.

During their analysis, exploiters often consult a variety of databases and other services to learn context and maintain situational awareness, as shown by the blue dotted line. Several resources, such as an operational picture display, a database, and an IRC⁶⁵ system, are kept in the background for situational awareness, which we show as optional conceptual flow. Sometimes, an analyst in the investigator role (e.g., MSA) assists in this process. Finally, when the exploiter completes their work—which generally consists of detecting and identifying objects, persons, or signals of interest—they record this information in some form of product. They then email this product to whoever initially requested the collection and store it in various product repositories. We have cross-hatched this part of the process as part exploiter, part reporter because different INTs do this differently depending on how standardized this process is.

Depending on the INT, the reporting process may involve additional steps to create more-complex reports and products. In many cases, a reporter helps compile additional information. With FMV, for example, typically one exploiter watches the video, another maintains communication with the mission-control element, and a reporter compiles information on significant events during the shift to enable them to write the final mission report.

Finally, at the end of the shift, the PED crew issues various reports. The supervisor provides quality control, records statistics for mission assessment processes, and prepares the crew for the changeover to the next shift. Sometimes, these steps are performed piecemeal throughout the shift and compiled at the end. The free-floating “Mission Management” text on the far left of the figure indicates that the report subsequently enters various administrative and assessment processes that are outside the PED process but support it directly.

Although most discussion about making the AF DCGS more efficient centers on the exploiter role, the data flow map emphasizes that other roles—notably that of the reporter—can affect the overall process. Improvements in all areas will be needed if a PED crew is to be able to

⁶⁵ The most common IRC client is mIRC. Since its inception in 1994, mIRC has become so widely known that many refer to any IRC client as (somewhat redundantly) “mIRC chat.” According to its creator, Khaled Mardam-Bey, the *m* in the abbreviation has no confirmed meaning. As he puts it in his official frequently asked questions list: “It quite possibly stands for ‘moo,’ or perhaps even MU.” The word *MU* links to a page that indicates that it is a Chinese ideogram for “no-thing” (mIRC.com, “Personal FAQ,” webpage, 2020).

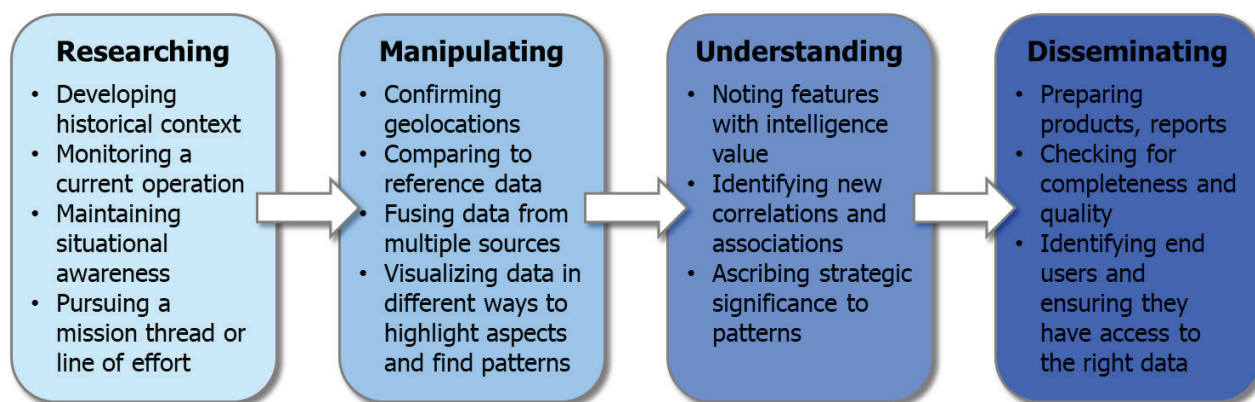
perform their tasks more rapidly. Otherwise, the bottleneck may just move to a different role or position.

DART

The workflow in the DART is not quite as amenable to a data flow map analysis because it differs considerably depending on the type of problem or data fusion being worked on. Moreover, because these activities are not usually tied to an active mission, there is not always a clear anchor point (indeed, sometimes the analysis can be self-initiated rather than in response to a request for information). However, we note that almost all of the DART's activity can be considered a form of the *investigator* role. We now take a closer look at this role and describe the four main types of analysis that DART investigators perform.

These four types of analysis, which are usually conducted more or less sequentially, are *researching*, *manipulating*, *understanding*, and *disseminating*. We give examples of these kinds of analysis in Figure 2.4. Note that these stages assume that the data in question have already been collected and processed. The data also might already have been exploited (the DART rarely performs first phase analysis). The analyst's objective throughout these stages is to complete and transmit information, usually through fusing information from multiple data sources, to answer a broad intelligence question.

Figure 2.4. DART Investigator Functions



Researching is usually the first stage of an investigation. This can include everything from attending PMBs (a relatively passive activity) to exploring existing reports and talking with others who may have looked at a similar topic or geographic area in the past. In addition to developing historical context, an analyst might be actively involved in monitoring an operation to develop understanding or checking CCMD, air component, or other lists of key types of intelligence questions or threads that could help an analyst determine when an observation is broadly significant. The greatest challenge with this work phase is knowing where to look and what for. Although geographic area provides one foundation by which to scope potential

research efforts, the most relevant previously exploited or analyzed information must be discoverable, and the analyst must be able to distinguish significant insights.

Manipulating starts when an analyst or team receives processed data and/or previously exploited information and begins doing something new with it. The greatest challenge with this work phase is managing the number of disparate programs an analyst must use. Much of this work involves comparing the particular collection being analyzed (e.g., the image) to reference data to find anomalies or fusing data from multiple sources to seek new patterns. Using different visualization tools is also an important part of the manipulation process for human analysts.

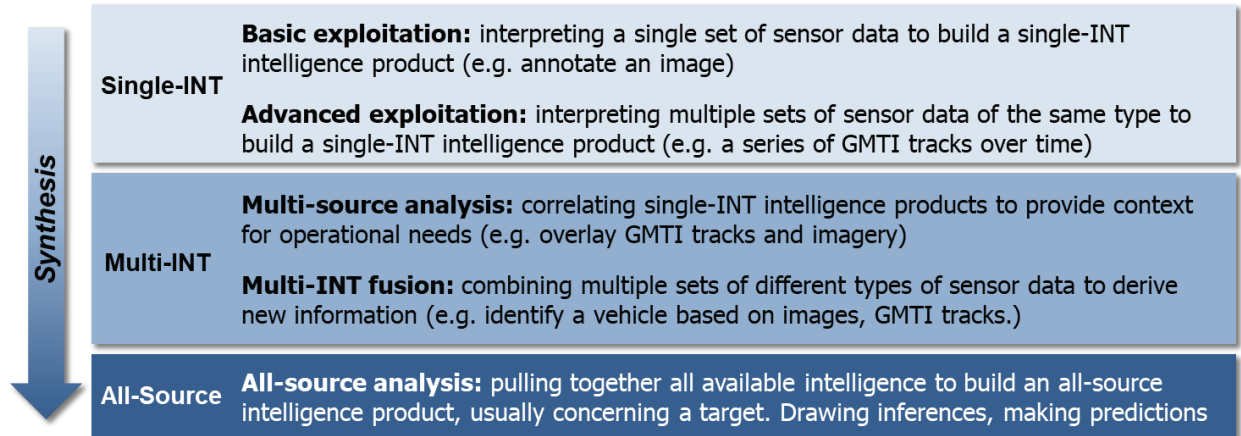
Understanding is the process by which an analyst or team determines the significance of observations (or lack thereof). The greatest challenge with this work phase is acquiring sufficient context to make sound judgments.

Disseminating involves making sure the right people have access to the right information at the right time. Sometimes, the result is a PowerPoint presentation. Other times, it may be a posting to an intelligence site or even a blog. The greatest challenge with this work phase is ensuring that all the formatting and legal requirements have been met and that the work is received at the correct locations.

Data Flow and the Spectrum of Synthesis

Another way to think about data flow in relation to the types of information used is the “spectrum of synthesis” shown in Figure 2.5. The spectrum of synthesis shows how data flow in a very general sense through the analysis process. First, there are single-INT processes of basic and advanced exploitation. The basic process involves a single set of data. Advanced processes might involve comparing several sets of data, such as looking at two different images of the same site. This is generally performed by the PED crews on the ops floor today. The DART is responsible for multisource analysis and multi-INT fusion, although some multisource analysis is part of the investigator role in PED crews as well. The distinction between multisource and multi-INT analysis was made specifically to distinguish between those roles. In multisource analysis, products are pulled together and correlated, but true multi-INT analysis can go deeper and involve additional work with the raw data. The end of the spectrum is all-source analysis, which we use broadly here to include target system analysis (TSA) and other functions that are not done in real time.

Figure 2.5. Spectrum of Synthesis



SOURCE: Adapted from Alkire et al., 2016, p. 46.

NOTE: GMTI = ground moving target indicator.

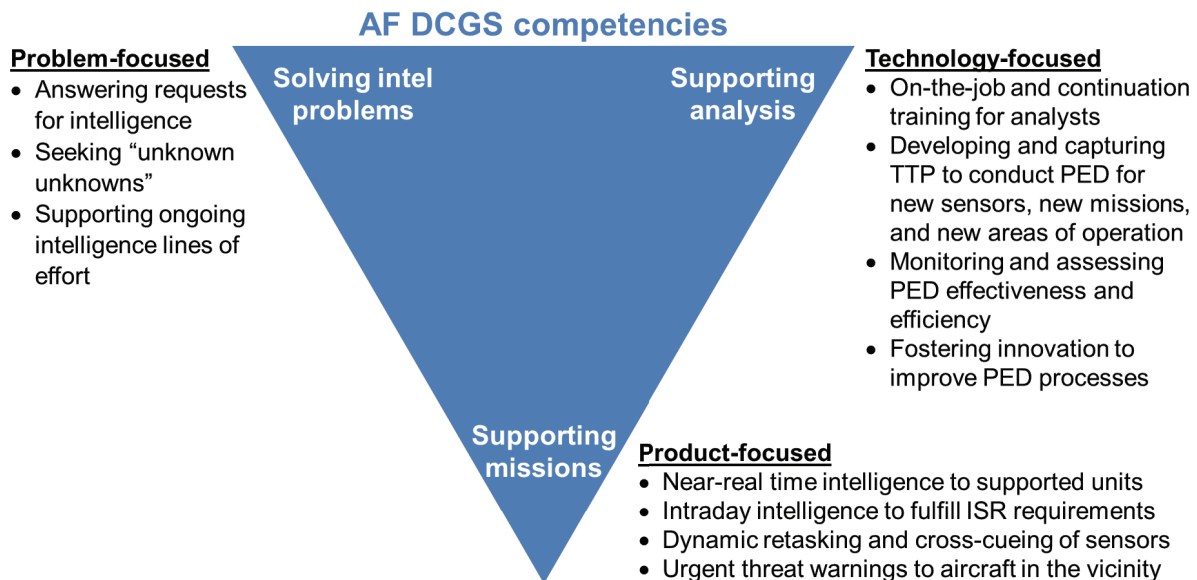
As discussed in Volume 1 and later chapters of this report, AI/ML may be able to push human effort down the spectrum of synthesis toward more-challenging work. But it is important to realize that synthesis is a chain that must be supported all the way down, or one risks creating unintended bottlenecks in the PED process.

AF DCGS Competencies

As the overview of organization, roles, and data flows just presented suggests, the AF DCGS is responsible for many different types of activities that support warfighters and advance national security objectives in various ways. We conclude this overview by thinking through a business model for the AF DCGS that will support technological change today and lay the groundwork for more change in the future.⁶⁶ A key element is the value proposition: the unique mix of products and services that an organization offers. There are three general competencies that the AF DCGS provides to the warfighter that we (loosely) call *supporting missions*, *supporting analysis*, and *solving intelligence problems*. These are summarized in Figure 2.6. These competencies are usually provided by analysts in different roles. This chart represents our analysis and observation of AF DCGS training documents and operations.

⁶⁶ Of course, the AF DCGS is not a business, but considering such organizational principles as this can be an instructive exercise.

Figure 2.6. AF DCGS Competencies



The AF DCGS supports missions by (1) providing near–real-time intelligence to supported units on the ground; (2) providing intraday intelligence (within approximately 24 hours) to fulfill ISR requirements (e.g., with high-altitude imagery); (3) conducting dynamic retasking and cross-cueing of sensors; and (4) providing urgent threat warnings to ISR platforms and other aircraft in the vicinity. The threat warning aspect of AF DCGS is particularly salient because it is one of the few Air Force–centric missions that the AF DCGS regularly conducts today (most AF DCGS PED supports other components). Although PED consists largely of single-INT processes, multi-INT fusion and cross-cueing are increasingly important as AF DCGS “customers” come to expect more. Today, this competency is generally fulfilled by airmen performing the *exploiter* and *reporter* roles on the ops floor. As discussed in Chapter 6, AI/ML might be able to perform or assist with many of the tasks in this competency, freeing analysts to focus on other competencies.

Supporting analysis consists of five interrelated functions: (1) provide on-the-job and continuation training for analysts; (2) develop and capture TTPs to conduct PED for new sensors, new missions, and new areas of operation; (3) monitor the effectiveness and efficiency of analytic processes (e.g., conducting assessments); and (4) foster innovation to improve these processes, which may be based on these assessments or achieved through other means. These functions are done by various analysts and groups in the *supervisor* role. As we discuss later, this competency may be expanded to include helping to monitor, train, and deploy future AI/ML applications.

Solving intelligence problems tailored to user needs is both a link to the history of the AF DCGS in supporting ground operations in Afghanistan and Iraq and a link to its future. The AF DCGS increasingly answers requests for information regarding broader intelligence problem sets

that bear significance beyond the air tasking order cycle or mission day. This has partly grown out of the regional alignment of DGS sites with other CCMDs, such as the U.S. Indo-Pacific Command, where intelligence needs are less driven by COIN/CT operations than they are in U.S. Central Command. Post-ingestion or “forensic” analysis of WAMI data could also be part of this process if it is tied to a request for information. Developing this type of understanding is inherently a longer-term process of knitting together data and information. These functions are primarily performed today by analysts in the *investigator* role in the DART.

Although the AF DCGS today works in all of the above competencies to some degree, the general weight of effort is on supporting missions. This is partly because of the challenges discussed in Volume 1: High demand for short-term PED, burdensome reporting requirements, and other factors mean that significant analyst attention is consumed with these tasks. The ops floor reflects this emphasis, with crews largely organized around platforms and INTs. As we discuss in later chapters, the use of AI/ML for these tasks might allow the AF DCGS to rethink how it organizes and allocates the weight of effort for human analysts across the three competencies, with potential benefits for addressing the enduring challenges of effectiveness, efficiency, use of human capital, and agility.

3. Improving Efficiency, Effectiveness, Human Capital, and Agility: Lessons from Historical Case Studies

Volume 1 identifies efficiency, effectiveness, the use of human capital, and agility as major objectives for improving the AF DCGS.⁶⁷ In this chapter, we take a deeper look at the challenges involved in meeting these objectives as illustrated by past operations and experiences. We select four case studies that highlight the challenges associated with each of the above objectives: dynamic targeting (efficiency), deliberate targeting (effectiveness), airman resiliency (use of human capital), and humanitarian assistance and disaster relief (agility).⁶⁸

Dynamic Targeting

Background

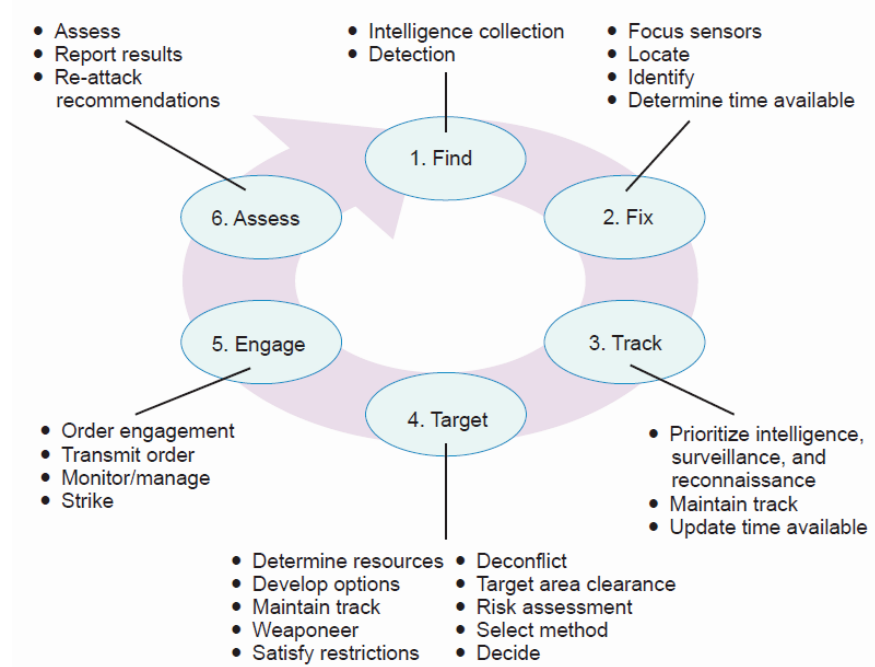
Dynamic targeting is the process of prosecuting targets of opportunity—both unplanned and unanticipated.⁶⁹ These may include both fixed and mobile targets. Joint doctrine defines the standard process for dynamic targeting in terms of six largely sequential steps: *find*, *fix*, *track*, *target*, *engage*, and *assess* (F2T2EA), as shown in Figure 3.1.

⁶⁷ Menthe et al., 2021.

⁶⁸ We look primarily at examples beyond recent COIN/CT operations. Although ISR support to such missions as understanding civilian patterns of life and striking a high-value individual are among the most common in today's fight, the AF DCGS must prepare for other missions as well.

⁶⁹ There are two types of targeting: deliberate and dynamic. Notably, as indicated in joint doctrine: "Neither is indicative of the target to be engaged but is more closely aligned with the planning phase in which the target is identified and prosecuted." In terms of dynamic targeting, an *unplanned target* is a known target that was not nominated for engagement because of priority or some other consideration, while an *unanticipated target* is an unknown target or one that was not expected to be in the area (Joint Publication 3-60, *Joint Targeting*, Washington, D.C., January 31, 2013, p. II-1).

Figure 3.1. Find, Fix, Track, Target, Engage, and Assess Cycle



SOURCE: Joint Publication 3-60, 2013, p. II-23.

The first step is to *find* the target. Initial analysis of the adversary might indicate where the target is likely to be located and what operating procedures the target might employ, which provides a starting point for employing ISR assets. Collection managers task these assets to detect emerging targets and characterize them (e.g., determine whether they need to be engaged), shaping the tasking instructions using the initial intelligence.

Once a target has been found, it must be *fixed*. Fixing is the process of identifying the target and its location with sufficient accuracy to employ weapons if needed. This step may require additional collection. Fixing the target generally requires the use of precision point mensuration to provide geographic coordinates in all three dimensions. If possible, the time available to engage the target may be estimated.

After the target has been fixed, it must be *tracked* so that operators will know if it moves before the F2T2EA process can be completed. Although the target is being tracked, other targeting functions, such as weaponearing,⁷⁰ point mensuration, and collateral damage estimation may be conducted to support the attack.

⁷⁰ *Weaponearing* means choosing the right weapon for the target. Formally, it is “the process of determining the quantity of a specific type of lethal or nonlethal means required to create a desired effect on a given target” (Joint Publication 3-60, 2013, p. II-15).

The next step is to *target* the target.⁷¹ A host of related targeting functions occur during this step as appropriate. These include estimating collateral damage, planning for poststrike assessment, and receiving approval to strike.

The final steps are to *engage* the target and then *assess* the effectiveness of the engagement. Often, this battle damage assessment requires additional collections to confirm the weapon effects.

Mobile targets are perhaps the most challenging type of dynamic target because they can employ their mobility as a survivability measure. The key to defeating this kind of target with the F2T2EA process is speed: completing the dynamic targeting process swiftly enough to engage the target before it can move (e.g., to a concealed or hardened position where it may be unfeasible to find or strike). In some situations, the platforms employed in the strike may have fuel or route requirements, adding to the planning constraints. This drives a need for tight tactical and operational coordination.

In target-poor environments, there may be only a few mobile targets to process at a time. But in the opening days of a more conventional fight, the number of mobile targets can rise rapidly. Improvements in AF DCGS analytic capabilities may result in “finding” more targets initially, but if the remaining steps are not optimized as well, many of these targets may disappear before they can be fully prosecuted. This difficulty was seen very clearly in the Scud missile-hunting operations during the Persian Gulf War.

Case Study

Stopping Saddam Hussein’s mobile Scud missile attacks on Israel was an early priority in Operation Desert Storm (1991) because the prospect of Israeli involvement threatened to crack the fragile coalition of Western and Arab states.⁷² To persuade Israel to remain on the sidelines, the United States made the destruction of Iraqi Scud missiles a major objective of the air campaign.⁷³ Therefore, a total of 4,754 anti-Scud sorties were planned during the 43-day air campaign in the Kuwaiti theater.⁷⁴

The Iraqi military had two primary means of launching Scud missiles: fixed launchers and mobile transporter erector launchers (TELs). Initial intelligence estimates put the Iraqi Scud inventory at 28 fixed launchers and 36 TELs (although post-war analyses later labeled the TEL

⁷¹ Although perhaps redundant, this usage is standard.

⁷² William Rosenau, *Special Operations Forces and Elusive Enemy Ground Targets: Lessons from Vietnam and the Persian Gulf War*, Santa Monica, Calif.: RAND Corporation, MR-1408-AF, 2001, p. 43.

⁷³ J. Michael Kennedy, “U.S. Rushes Defenses to Israel: American Troops to Operate Two Patriot Batteries: Gulf War: The Move May Forestall an Immediate Retaliatory Strike by the Jewish State After Two Attacks on Tel Aviv. Allied Warplanes Step Up the Search for Scud Mobile Launchers,” *Los Angeles Times*, January 20, 1991.

⁷⁴ Mark David Mandeles, Thomas Hone, and Sanford S. Terry, *Managing “Command and Control” in the Persian Gulf War*, Westport, Conn.: Praeger Publishers, 1996, p. 76.

estimates “sketchy and too low”⁷⁵). The first three days of the air campaign called for 356 anti-Scud sorties against well-known Scud sites, plus Scud manufacturing and maintenance sites to cripple Iraq’s capabilities. Unfortunately, the fixed launcher sites proved to essentially be decoys, diverting attention away from the dispersed, mobile TELs that launched 12 missiles over those first three days.⁷⁶ Having failed to suppress Iraq’s missile-launching capability by destroying fixed sites, military planners set their sights on the mobile Scud launchers.

The goal was to find and destroy the TELs before they could fire. However, poor weather, the size of the area in which the TELs could hide, and Iraqi TTPs made them elusive targets: In addition to effective concealment prior to launch, Iraq cut launch preparation times down to just 30 minutes.⁷⁷ This deception and speed made executing the F2T2EA process prelaunch nearly impossible.

Unable to strike Scud launchers prior to missile launch, U.S. and coalition forces adapted their TTP to attempt to complete the F2T2EA process immediately following a launch. To do so, planners established “kill boxes” in areas of suspected Scud activity. Strike and other support aircraft patrolled the kill boxes seeking to detect a launch, identify the firing location, acquire the TEL with onboard sensors, and deliver ordnance effectively before the Scud missile crew could pack up the TEL and slip away. Unfortunately, “by wars end, nearly every type of strike and reconnaissance aircraft employed in the war participated in the attempt to bring this threat under control, but with scant evidence of success.”⁷⁸

Finding a TEL after it had launched a missile was the easy part: “once ignited, Scud missile motors produced a visible and very hot plume, and this plume was something that pilots overhead could see and satellites in space could detect.”⁷⁹ The problem lay with the subsequent steps in the F2T2EA process. On-board sensors of patrolling aircraft were unable to fix mobile Scud launchers in a timely manner because the launchers’ infrared and radar signatures were “virtually indistinguishable from trucks and other electromagnetic ‘clutter’ in the Iraqi desert.”⁸⁰ The identification process was confounded mainly by commercial tractor-trailer trucks and decoys. One statistic illustrates the challenge: Of the 42 cases of visual observation by patrolling strike aircraft, only eight aircrews could sufficiently acquire the TELs to deliver ordnance. And

⁷⁵ This was the exact phrase from DoD’s final report to Congress (DoD, *Conduct of the Persian Gulf War Final Report to Congress*, U.S. Government Printing Office: Washington, D.C., 1992, p. 150.

⁷⁶ Thomas A. Keaney and Eliot A. Cohen, *Gulf War Air Power Survey: Summary Report*, Washington, D.C.: U.S. Government Printing Office, 1993, p. 86.

⁷⁷ Perry Jamieson, *Lucrative Targets: The U.S. Air Force in the Kuwaiti Theatre of Operations*, Washington, D.C.: Air Force History and Museums Program, 2001, p. 48.

⁷⁸ Keaney and Cohen, 1993, p. 17.

⁷⁹ Mandeles, Hone, and Terry, 1996, pp. 74–76.

⁸⁰ Rosenau, 2001, p. 34.

even in those instances, ineffective battle damage assessment left coalition forces unclear on how many mobile Scuds were destroyed.⁸¹

Various ISR aircraft and satellites were employed to overcome the limitations of the on-board sensors of strike aircraft, but the PED process often was too slow to make a significant difference. For example, the joint surveillance and target attack radar system (JSTARS) was employed to provide additional data, but it did not lead to greater success in destroying mobile Scud launchers “because the time it took to analyze JSTARS images and then communicate with patrolling F-15Es proved too long.”⁸²

U.S. Space Command’s Scud-warning process was faster—using satellite observations, it could provide an approximately four-minute warning to U.S. Central Command’s command-and-control centers and Patriot battery commanders. This allowed Patriot battery commanders to effectively engage some of the incoming missiles, but dissemination of the warnings was either too fractured or too slow to effectively cue patrolling strikers to the mobile Scud TELs.⁸³ In the aftermath of the Gulf War, the Defense Science Board concluded that “a capability to find and destroy . . . Scuds before they launch implies hitherto unachieved integration and a new level of processing of surveillance data.”⁸⁴

Lessons

This case highlights the importance of *efficiency* in the analysis process, and specifically, the need for timeliness. It also highlights another issue: Although an efficient process clearly is essential for dynamic targeting, an efficient process is not always sufficient to improve operational outcomes. The entire F2T2EA chain must be timely, and analysis is but one part of that longer process.⁸⁵

Much has been done to improve PED efficiency since Operation Desert Storm, but expectations for PED efficiency have also risen. The ability to *supply* near-real-time intelligence has engendered a *demand* for near-real-time intelligence: Process timelines that would have been considered efficient in 1991 would be considered broken today. Moreover, improved adversary TTPs will continue to force a tighter analytic cycle in future operations. As discussed in Volume 1, making analytic processes more efficient to meet these challenges and demands is one area where AI/ML could make a significant contribution.

⁸¹ Keaney and Cohen, 1993, p. 87.

⁸² Mandeles, Hone, and Terry, 1996, p. 78.

⁸³ Mandeles, Hone, and Terry, 1996, p. 77.

⁸⁴ Defense Science Board, *Lessons Learned During Operations Desert Shield and Desert Storm*, Washington, D.C., June 8, 1992, p. 74.

⁸⁵ This fact cautions against using operational outcomes as the sole measure of PED value because it may easily lead to situations in which the PED process was handled very well, but the desired outcome was not achieved for other reasons. It also leads us to note areas outside PED where inefficiencies must be addressed.

Deliberate Targeting

Background

Deliberate targeting—which is related to the concept of strategic attack—dates to the earliest airpower theorists. Although there are variations between the ideas of different theorists, the general concept is to strike key strategic targets in a way that creates outsized effects across the opposing military. The targets usually involve key functions, such as early warning systems, command and control, and communications. Targeteers conduct a TSA to determine which targets are most critical. If the TSA is successful, the analyst might uncover targets that, if struck, could have a crippling effect on the adversary.

Conducting a proper TSA requires multiple techniques and is often quite data intensive. A TSA report can take months to produce. Where dynamic targeting requires rapid analysis and effective operational synchronization, the intelligence work required for supporting this kind of deliberate targeting has historically been quite different. It often requires collecting large amounts of data over time, patiently waiting for intelligence gaps to be closed, accessing the information, wading through that information, and processing it in a way that yields new understanding of the adversary's strengths and weaknesses.

Case Study

Operation Allied Force (OAF) illustrates some of the challenges of providing ISR support to a strategic attack. The stated objectives for the United States and NATO in OAF were to (1) ensure a verifiable stop to all military action and the immediate ending of violence and repression; (2) ensure the withdrawal from Kosovo of the military, police, and paramilitary forces; (3) agree to the stationing of an international military presence in Kosovo; (4) agree to the unconditional and safe return of all refugees and displaced persons and unhindered access to them by humanitarian aid organizations; and (5) provide credible assurance that Serbian President Slobodan Milošević would negotiate a political settlement for Kosovo in conformity with the Rambouillet Accords and the Charter of the United Nations.⁸⁶

To achieve these objectives, NATO launched an air campaign beginning on March 24, 1999, and ending 78 days later, when Milošević accepted NATO's conditions for surrender.⁸⁷ Although NATO was ultimately successful in achieving its overall objectives, a lack of clarity in targeting strategy led to a suboptimal application of air power and ISR assets.

Supreme Allied Commander in Europe GEN Wesley Clark initially planned simply to “grind away” at Serbian forces.⁸⁸ The goal was to “degrade” and “damage” Yugoslavian assets until

⁸⁶ NATO, “The Aims of the Air Campaign,” October 30, 2000.

⁸⁷ The terms were essentially the same as the stated objectives for OAF (CNN, “Milosevic Accepts Peace Plan, Finnish Envoy Says,” June 3, 1999).

⁸⁸ Eric Schmitt, “Weak Serb Defense Puzzles NATO,” *New York Times*, March 26, 1999.

Milošević accepted NATO terms. The plan was “driven by the assumption that the operation would entail, at most, a two- to three-day series of air strikes directed at approximately 50 targets.”⁸⁹ NATO’s first strikes were mainly against “enabling targets,” such as anti-aircraft artillery that, once eliminated, would allow for easier strikes against other assets, but they also targeted critical infrastructure (e.g., the electrical power grid of Kosovo’s capital, Pristina) and ground force assets (e.g., Serbian army barracks and supply depots).⁹⁰ Unfortunately, these initial planned attacks did not deter the Serbs, who pressed ahead with their plan to drive as many ethnic Albanians as possible out of Kosovo to expose the Kosovo Liberation Army, which was believed to be intermingled with the civilian population.

When it became clear that the air campaign would not end quickly, General Clark pressed his staff to identify 5,000 additional target candidates (although he later lowered the goal to 2,000 target candidates when his staff convinced him the larger figure was unrealistic). Targeteers scrambled to come up with additional candidates to meet the goal, which proved to be a considerable challenge. One analyst labeled the NATO alliance’s target development a “mechanical process of meticulous selection with little true military justification.”⁹¹

The NATO campaign ultimately succeeded, but many analysts subsequently contended that it was the political pressure created by the bombing campaign rather than the destruction of any specific targets that had the most effect. As one analyst put it: “It was not *what* we bombed, but *that* we bombed.”⁹²

Lessons

This case highlights the importance of *effectiveness* in intelligence analysis. In this case, the operation was successful, but the scramble to find targets in OAF demonstrated to many that strategic planning processes for deliberate targeting needed to be improved. And later conflicts continue to indicate room for improvement. Operation Odyssey Dawn (2011) revealed that some of the same problems from the earlier era persisted:

Specifically, since the late 1990s, the combination of force restructuring, operational needs in a counterinsurgency environment, and service and DoD efficiency initiatives contributed to the atrophy of targeting capabilities across the board During that same time period, technological advances and new platforms, sensors, and munitions similarly transformed targeting requirements—the classic targeting folders and weaponeering process had changed into something both digital and dynamic. The result, underscored by experiences in Odyssey Dawn (the operation to enforce United Nations Security Council

⁸⁹ Lambeth, 2001, p. 199.

⁹⁰ Lambeth, 2001, p. 22.

⁹¹ William M. Arkin, “Smart Bombs, Dumb Targeting?” *Bulletin of the Atomic Scientists*, Vol. 56, No. 3, May/June 2000, p. 48.

⁹² Lambeth, 2001, p. 86.

resolution 1973 in Libya), is that Air Force targeting now lacks sufficient capacity to remain effective within the context of future planning scenarios.⁹³

Currently, efforts are underway in the 363rd ISRW to integrate TSA into the 72-hour air tasking order cycle. A key principle is the recognition that targets are not idle when attacked. The adversary makes combat repairs, opens new facilities to recover losses, and employs creative workarounds. Furthermore, tactical, operational, and strategic objectives of friendly forces might shift as the conflict unfolds. As a result, there is a need for continued TSA throughout active combat operations to update deliberate targeting priorities as the conflict progresses.

TSA analysts spend significant time tagging images with basic information today, which adds to the laborious nature of that process. This is necessary for two reasons: first, because the annotated image product sometimes cannot be located without the tag, and second, because the tagging during the initial PED process may have been efficient but was not *complete*. Although the requested EEIs accompanying the original tasking order may have been satisfied, not everything that could be gleaned from the image was recorded. Just as PED is but one part of the F2T2EA chain, it can also be the first stage in a longer analysis process. The effectiveness of earlier phases of analysis has implications for subsequent steps. As discussed in Volume 1, AI/ML (along with other, less sophisticated automation methods) may be able to improve AF DCGS operational effectiveness, both by providing data-management options to make information more available and by performing some of the more routine tasks that prevent analysts from focusing on strategic intelligence problems.

Airman Resiliency

Background

A previous PAF study that surveyed AF DCGS analysts confirmed two distinct types of mental workload concerns: the tedious, “mind-numbing” tasks associated with low-tempo operations—often from watching video where nothing was happening—and, at the other end of the spectrum, frantic efforts associated with high-tempo operations.⁹⁴ This matches well the long-standing recognition in the organizational psychology literature that both underload and overload can be drivers of occupational burnout.⁹⁵ A separate 2012 study found that AF DCGS analysts reported unusually high levels of stress and exhaustion compared with nonintelligence

⁹³ Kimminau, 2012, pp. 124–125.

⁹⁴ Menthe et al., 2015b.

⁹⁵ *Underload* is defined as “tedium and monotony,” whereas *overload* is “too many demands with too few resources” (Christina Maslach, Wilmar B. Schaufeli, and Michael P. Leiter, “Job Burnout,” *Annual Review of Psychology*, Vol. 52, February 2001, pp. 405, 414).

analysts.⁹⁶ As yet another study explained, the environment at the time was stressful because of the effects of the near-continuous surge conditions under which the AF DCGS had been operating:

The most problematic stressors among DCGS intelligence operators continue to be operational in nature, such as high workload and not enough manpower to accomplish all tasks, organizational conflict associated with being assigned multiple tasks competing for time, interpersonal conflict, and shift work.⁹⁷

These stressors incur costs to the organization and interfere with training. As Gen Herbert J. Carlisle put it, “One of the challenges we’re having with respect to the RPA enterprise is we can’t get breathing room to do anything.” He explained that crews “don’t get to practice their entire mission sets” and were doing “zero continuation training because they’re all engaged in the fight.”⁹⁸ And the concerns go beyond work-related issues. They can also threaten mental health. As one study found:

Enlisted RPA intelligence specialists displayed significantly higher incidence rates for substance abuse/dependence, family circumstance problems, and maltreatment related mental health categories, and for all mental health outcomes combined . . . [they] also displayed statistically higher incidence rates for life circumstance problems and posttraumatic stress disorder. . . . Military policymakers and clinicians should recognize that RPA intelligence personnel have increased mental health risk while performing their duties.⁹⁹

Case Study

To address these growing concerns, in 2015, the Air Force began embedding small groups of medical and psychological professionals at several AF DCGS sites.¹⁰⁰ These Airman Resiliency Teams (ARTs) provide a variety of mental health services, as well as ways to improve the work environment. The ARTs are credited with helping bring significant reductions in the suicide rate. The Air Force Medical Service reported:

Integration of the 480th ISR Wing ART took place in 2015. The embedded teams have made a huge difference in the physical and mental wellbeing of its Airmen.

⁹⁶ For example, 29 percent of intelligence analysts reported serious exhaustion at least once per week, versus only 6 percent in other fields (John K. Langley, *Occupational Burnout and Retention of Air Force Distributed Common Ground System [DCGS] Intelligence Personnel*, dissertation, Santa Monica, Calif.: Pardee RAND Graduate School, 2012).

⁹⁷ Lillian Prince, Wayne L. Chappelle, Kent D. McDonald, Tanya Goodman, Sara Cowper, and William Thompson, “Reassessment of Psychological Distress and Post-Traumatic Stress Disorder in United States Air Force Distributed Common Ground System Operators,” *Military Medicine*, Vol. 180, No. 3 Supplement, March 2015, p. 176.

⁹⁸ Gen Herbert “Hawk” Carlisle, former ACC commander, quoted in Mark Pomerleau, 2015.

⁹⁹ Kris Anthony Ostrowski, “Psychological Health Outcomes Within USAF Remotely Piloted Aircraft Support Career Fields,” dissertation, Daytona Beach, Fla.: Embry-Riddle Aeronautical University, June 2016, p. iv.

¹⁰⁰ Peter Holstein, “Airmen Resiliency Team Provides 480th ISR Wing with Medical, Psychological and Spiritual Care,” Air Force Medical Service webpage, Surgeon General Office of Public Affairs, May 24, 2017.

Long considered a career field with higher risk of suicide, 2016 saw zero suicides for the unit.¹⁰¹

In addition, the ARTs looked to work-related improvements. They found that AF DCGS airmen were becoming “more tired, easily distracted and less effective toward the end of shifts lasting longer than eight hours.”¹⁰² A supporting study found that retention among analysts was also a concern because shifts were often 14 hours long: “Work schedule was the factor most listed as affecting the decision to re-enlist.”¹⁰³

Implementing the recommended change to shorter shifts had many measurable benefits. In terms of quality of life, airmen reported improvements on every measure, including not feeling exhausted.¹⁰⁴ There were also product improvements in terms of effectiveness. The error rate for high-altitude imagery and FMV in their primary area of responsibility was 12 percent and 8 percent lower, respectively, and airmen showed “improved initiative” and were less likely to fail to use additional resources in their work.¹⁰⁵

Lessons

ARTs show the importance of protecting *human capital* in the AF DCGS. Caring for human capital is not just good for the humans but also is good for the organization. Therefore, it is important to be mindful of how automation of analytic processes within the AF DCGS will affect the texture of the work done and whether that will lead to underload or overload. Volume 1 suggests that AI/ML, by taking over the more tedious tasks, can free analysts for jobs that are both more professionally satisfying and make better use of human capabilities. At the same time, the AF DCGS should be careful to introduce new tools in a way that helps analysts rather than creates new burdens, as discussed in Chapter 8 of this report.

Humanitarian Assistance and Disaster Relief

Background

Periodically, the Air Force is called on to conduct or support humanitarian assistance and disaster response (HA/DR) operations, typically in response to large-scale events that overwhelm local authorities. In these situations, time is often an enemy: Lives may still be in danger as a

¹⁰¹ Holstein, 2017.

¹⁰² Holstein, 2017.

¹⁰³ Although shifts were nominally 12 hours long, they were actually closer to 14 hours, including the briefings and changeover tasks on either end (548th Operational Support Squadron, *Sustainable DCGS: DGS-2 Pilot Study 18 Mar–25 Jun 2015*, Beale Air Force Base, Calif., July 2015, p. 9).

¹⁰⁴ Jeremy Didier, “Sustainable DCGS: DGS-2 Pilot Study 18 Mar – 25 Jun 2015,” 548th Operational Support Squadron, July 2015.

¹⁰⁵ Didier, 2015.

result of the direct effects of the disaster, and more may perish if the required food, water, medicine, and other resources cannot be delivered in a timely manner. Understanding the situation on the ground is critical to delivering the right aid to the right place at the right time.

In some situations, the Air Force ISR community is asked to provide some of that situational understanding. The greatest challenge here often lies in the dissemination of information. The AF DCGS has provided exploitation for HA/DR operations via special products cleared for release.¹⁰⁶ But the Air Force ISR community was built and optimized to provide operational warfighters with the information they need to succeed against a foreign adversary, and commanders typically have the clearances required to see this information.¹⁰⁷ In a HA/DR situation, however, decisionmakers might not have the right clearances—indeed, in situations on foreign soil, the most-important decisionmakers might not even be U.S. citizens.¹⁰⁸

For U.S. operations, the Air Force ISR community must work through intelligence oversight rules to ensure that appropriate Proper Use Memorandums are in place. In both domestic and foreign situations, the information must be made releasable to appropriate authorities, either through declassification or working through the Foreign Disclosure process. All of this adds time, which can be detrimental to operational success.

Case Study

The Air Force played a vital role in the disaster response to Hurricane Katrina. A total of 13 airborne ISR assets were employed by six different government agencies—Air Force, ANG, Civil Air Patrol, Navy, Defense Intelligence Agency, and Customs and Border Protection—to assess damage and facilitate relief and recovery. In total, these assets flew 361 sorties, during which they collected 1,128 images and recorded 117 hours of FMV. These data were analyzed across six exploitation nodes and disseminated to 22 customers throughout the affected area.¹⁰⁹

Working through intelligence oversight issues and “sanitizing” the data to produce an unclassified product was necessary but time consuming. In an interview a year after the disaster, the Assistant Secretary of Defense for Homeland Defense, Paul McHale, noted that the organizations involved often had different communications equipment, security protocols, and standards. Although McHale was confident that top-level decisionmakers were sufficiently connected, he felt that rapid deployment and emergency assistance capabilities still need to be shored up. He concluded:

¹⁰⁶ U.S. Air Force, 2015b.

¹⁰⁷ This is not always true for all types of intelligence, unfortunately, but those issues go beyond the scope of this project.

¹⁰⁸ The information also must be promptly transmitted to any ground units directly assisting the affected population.

¹⁰⁹ U.S. Air Force, “Air Force Support to Hurricane Katrina/Rita Relief Operations: By the Numbers,” Washington, D.C.: Headquarters U.S. Air Force, October 2005, pp. 13–14.

We in DoD have a duty to work with our interagency partners in order to ensure that civilian capabilities are properly planned, effectively resourced, and are well coordinated with DoD to ensure that once we get downrange, our national response will achieve unity of effort.¹¹⁰

The U.S. disaster response to the Great East Japan Earthquake of March 2011, Operation Tomodachi, offers another example of how ISR assets can be effectively used in HA/DR but further underscores the longer timelines that can arise when attempting to disseminate ISR products to foreign decisionmakers. In this case, all four services, along with numerous governmental agencies, such as the U.S. Department of Energy (DoE) and the Nuclear Regulatory Commission, worked with Japanese governmental and nongovernmental agencies in this massive effort. At the peak of Operation Tomodachi, the United States had deployed approximately 24,000 personnel, 189 aircraft, and 24 ships to the affected area. Japan also deployed more than 100,000 disaster relief personnel, 500 fixed-wing and rotary aircraft, and 60 ships.¹¹¹

ISR for the effort was primarily provided by Air Force RQ-4 Global Hawks and U-2s. Global Hawks were used to search for and photograph survivors, living areas, and infrastructure to assess damage, identify usable roads, and help Japanese officials decide what areas need to be prioritized. Global Hawks were also used to monitor heat levels within the Fukushima Daiichi's reactors.¹¹² The Air Force also employed U-2s to capture broad-area images of affected areas, and the 9th Intelligence Squadron at DGS-2¹¹³ analyzed the photographs.¹¹⁴

The challenge was to ensure that all appropriate decisionmakers were receiving the right ISR information in a timely manner. Since the United States was playing a supporting role during the disaster, ISR products needed to be communicated to Japanese rather than U.S. government or military decisionmakers. Approximately ten days passed between the initial request and official approval of new information-sharing protocols. This was actually relatively rapid for this kind of dissemination: Situations involving authorization to share intelligence data in a way not previously done can take anywhere from one to four months. Although this timeline was sufficient for Operation Tomodachi, it illustrates the difficulties that can arise.

¹¹⁰ Merrick Krause and Jeffrey Smotherman, "An Interview with Assistant Secretary of Defense for Homeland Defense: Paul McHale," *Joint Forces Quarterly Forum*, No. 40, 2006, p. 15.

¹¹¹ Andrew Feickert and Emma Chanlett-Avery, *Japan 2011 Earthquake: U.S. Department of Defense (DoD) Response*, Washington, D.C.: Congressional Research Service, R41690, June 2, 2011, p. 1.

¹¹² Seth Robson, "Global Hawk Invaluable After Japan Disasters," *Stars and Stripes*, September 12, 2011; Tony Capaccio, "Northrop Drone Flies Over Japan Reactor to Record Data," *Bloomberg*, March 17, 2011.

¹¹³ Beale Air Force Base, "9th Intelligence Squadron," May 22, 2012.

¹¹⁴ Stacy Foster, "U-2 Reconnaissance Aircraft Deployed to Aid Japan Relief Efforts," 51st Fighter Wing Public Affairs, U.S. Air Force, March 13, 2011.

Lessons

These HA/DR examples highlight the importance of *agility* in the analytic process: the ability to quickly change processes to support different types of operations and to work with different partners. The inherent unpredictability of disasters also highlights the need for agility. The disclosure process for classified information is, and must be, deliberate and thorough because intelligence sources and methods can be highly sensitive. Moreover, establishing a new procedure for disseminating data often requires coordination from many top-level decisionmakers not only in DoD but at the U.S. Department of State and other agencies. But if these processes cannot be completed in a timely manner, they will not be helpful to the relief effort.

Although future platforms and collection technologies may provide the United States with a better understanding of disaster-affected areas, as the rate of data collection grows, so too does the risk that it will overwhelm not only the PED process but also intelligence oversight, declassification, and the foreign disclosure processes.

Recently, some steps have been taken to prepare for potential information-sharing situations. The 480th ISRW has worked with the CCMDs to create templates for every type of data and classification they currently process, and knowledgeable individuals have indicated their hope that this could enable timelines to improve from days to hours, if fully implemented. If preapprovals can be established or disclosure decisions can be safely pushed to lower levels for decisionmaking, this should allow for more-rapid dissemination in an emergency.

Automation of analytic processes may also speed the delivery of information once these approvals and decisions have been made, but it is worth noting that, in these HA/DR operations, trained imagery analysts were able to assist because they could quickly learn how to recognize new kinds of activity and objects in a new geographical area. As we discuss in Chapter 4, it is unlikely that AI/ML applications will be so flexible: It is difficult to extend their performance outside their training data set. This is one reason why we isolate “agility” as its own challenge—to ensure that this important character is not lost in the drive to streamline an organization that has, over the past 15 years, matured largely in the context of COIN/CT operations.

4. Artificial Intelligence and Machine Learning: A Primer for AF DCGS Analysts

The AF DCGS presently finds itself at a unique intersection of technical need and technical opportunity. The need is evident. Some of the computer systems still in use today at the AF DCGS are decades old,¹¹⁵ presenting today's young analysts with unfamiliar interfaces and requiring data to be hand-jammed from one system to another. Modern analysis tools, such as ArcGIS Pro, cannot run on these old systems.¹¹⁶ The open architecture DCGS program, which includes a long-overdue hardware and software refresh using Intel x86 and Microsoft Windows 10, is currently being tested and fielded across the AF DCGS enterprise, but, without additional funding, the rollout is not expected to be complete until 2021. As of August 2018, persistent data access and system integration problems prevented adoption at any DGS site.¹¹⁷

Meanwhile, over the last decade, the fields of AI/ML have experienced a renaissance as new life has been breathed into old algorithms—notably deep learning—and computing power has grown as its cost has decreased.¹¹⁸ This explosive growth in AI/ML has made advances toward addressing many analysis and categorization problems, such as machine vision and automated translation, that have obvious applications to the AF DCGS. Thus, the AF DCGS is ripe for technical overhaul at a time when useful AI/ML algorithms are becoming available for integration into the workflow and when moderation of the high “surge” conditions that were the norm during OEF and OIF affords the enterprise space for reflection and change.¹¹⁹

The Air Force is already moving in this direction:

“What we’re trying to do is set the conditions to build an AI-ready culture,” [Lt Gen John N. T. Shanahan] said. “It’s not easy. This is uncomfortable. It’s a very different way of thinking about problems than we’ve used in the past. But the attitude is out there. The younger people are more receptive to this and they’re

¹¹⁵ The authors saw some SUN workstations that were literally in use when we first examined this issue in 2009—when we first recommended new hardware.

¹¹⁶ Environmental Systems Research Institute, “ArcGIS Pro 2.2 System Requirements,” webpage, undated.

¹¹⁷ Among the issues cited were slow access times for imagery data, limited communications network access, and frequent crashes.

¹¹⁸ This has been widely observed. See, for example, Roger Parloff, “Why Deep Learning Is Suddenly Changing Your Life,” *Fortune*, September 28, 2016.

¹¹⁹ These surge conditions raised the maximum mission hours to 246 per month for a PED crewperson versus steady-state operations of 150 hours. See Table 4.1 in Air Force Intelligence, Surveillance, and Reconnaissance Agency Instruction 14-153, 2014.

ready to jump on board yesterday. They’ve been asking us: What took you so long?”¹²⁰

The purpose of this short tutorial on AI/ML concepts is to empower the reader to ask the right questions about these methods, understand the challenges involved with their implementation, and alert them to potential vulnerabilities. We believe that this level of knowledge is important to make informed investment decisions regarding AI/ML and to understand the recommendations presented in Volume 1.

Historical Overview

Today, *AI* can loosely be defined as the use of computers to carry out tasks that previously required human intelligence, but the term has evolved considerably over the field’s six-decade history. AI pioneer Marvin Minsky defined it in 1968 as “the science of making machines do things that would require intelligence if done by men.”¹²¹ Early AI researchers generally framed their goals in terms of replicating human cognition, as the predominant view was that machines could only evince intelligent behavior if they were also intelligent. But once computers began to succeed at such tasks, humans tended to decide that these tasks did not demand cognition after all. John McCarthy, credited by some with coining the term *AI*,¹²² dubbed this the *AI effect*: “As soon as it works,” he complained, “no one calls it AI anymore.”¹²³

In time, AI researchers increasingly sought to define their field in a more general way, sidestepping thorny philosophical problems about the nature of the human mind. A particularly influential framing held that the goal of AI instead should be merely the design of “rational agents,” things that act, “so as to achieve the best outcome or, when there is uncertainty, the best expected outcome.”¹²⁴

ML is not the same as AI, although they are often conflated in contemporary discourse. The term *machine learning* originated from AI—Arthur Samuel invented it to describe his landmark 1959 checkers program.¹²⁵ But researchers at the time were often less than eager to be associated with AI; some saw it as a controversial research program seeking to create “machines with

¹²⁰ Marcus Weisgerber, “The Pentagon’s New Artificial Intelligence Is Already Hunting Terrorists,” *Defense One*, December 21, 2017.

¹²¹ Marvin Minsky, ed., *Semantic Information Processing*, Cambridge, Mass.: MIT Press, 1968, p. v.

¹²² Sam Williams, *Arguing A.I.: The Battle for Twenty-First-Century Science*, New York: AtRandom, 2002, pp. 18–19.

¹²³ Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies*, New York: Oxford University Press, 2014, p. 13.

¹²⁴ Stuart Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach*, 2d ed., Upper Saddle River, N.J.: Prentice Hall, 2002, p. 4.

¹²⁵ Arthur Samuel, “Some Studies in Machine Learning Using the Game of Checkers,” *IBM Journal of Research and Development*, Vol. 3, No. 3, July 1959, pp. 210–229.

minds, in the full and literal sense.”¹²⁶ The kind of statistical methods that came to dominate ML were in the comparatively modest field then called “pattern recognition.”¹²⁷

In recent years, however, the term *ML* has increasingly subsumed both the label and the field of AI. This conflation obscures the fact that many ML applications, such as logistic regressions and clustering algorithms, are not really “artificial intelligence,” even under the most expansive definitions of the field. This misconception stems in part from the fact that learning techniques have come to predominate AI research, including in such uses as theorem-proving and machine-language translation that early AI researchers believed had nothing to do with learning.

The brisk progress in AI research over the past decade stems substantially from the sudden success of a subfield of ML known as deep learning (see discussion later in this chapter). Although deep learning was first conceptualized in the 1980s, it was too computationally intensive to be practical at that time. In the 2000s, however, deep learning began showing some promise in research settings and rapidly eclipsed previous techniques for many challenging AI/ML tasks. In the space of a few years in the mid-2010s, deep learning rendered prior approaches to machine vision and automated speech recognition obsolete.¹²⁸ In conjunction with reinforcement learning—another existing technique whose results had previously been underwhelming—deep learning also swept the field of game-playing, achieving superhuman performance at Go and other games.¹²⁹ It is important to note, however, that deep learning is not usually sufficient in and of itself in performing most cognitive tasks; it must be paired with other processes and methods.

Artificial general intelligence (AGI) is a still-hypothetical technology that would endow machines with some equivalent of humans’ flexible cognitive abilities. There is no generally accepted definition of this term, which is used by different authors to mean anything from systems that are merely more adaptable to novel problems than today’s “narrow” AI to machines that duplicate or exceed human performance on all cognitive tasks. Experts also do not concur about the kind of technology that would be used to implement AGI. Most researchers believe that simply scaling up current techniques will prove insufficient and that fundamental breakthroughs will be required. The applications we discuss in this report do not require AGI, because AGI remains abstract and theoretical, and less-powerful measures should suffice.

¹²⁶ John Haugeland, *Artificial Intelligence: The Very Idea*, Cambridge, Mass.: MIT Press, 1985, p. 2.

¹²⁷ Pat Langley, “The Changing Science of Machine Learning,” *Machine Learning*, Vol. 82, No. 3, March 2011, pp. 275–279.

¹²⁸ Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton, “ImageNet Classification with Deep Convolutional Neural Networks,” Fernando Pereira, Christopher J. C. Burges, Léon Bottou, and Kilian Q. Weinberger, eds., *Proceedings of the Advances in Neural Information Processing Systems 25*, 2012.

¹²⁹ David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot et al., “Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm,” *arXiv preprint arXiv:1712.01815*, 2017.

Glossary

Table 4.1 presents some common AI/ML terms that we use in the remainder of this chapter. They are shown in an order such that successive terms build on the previous ones. This is, of course, just a sample of a rich subject. We discuss these in more detail when we introduce these concepts in the AI/ML methods section.

Table 4.1. Glossary of AI/ML Terms

Term	Definition
Supervised learning	ML that learns to approximate a function by mapping between provided input-output pairs (i.e., labeled training data)
Unsupervised learning	ML that attempts to learn a function describing the structure of unlabeled training data
Reinforcement learning	ML that seeks to learn to maximize an expected reward or utility function by interacting with an environment over time
Neural network	A family of ML techniques that are loosely inspired by biological brains (also commonly called <i>artificial neural networks</i> to distinguish them from actual networks of linked biological neurons)
Node	The fundamental building block of a neural network, sometimes referred to as an <i>artificial neuron</i> or also a unit
Layer	The intermediate structure of a neural network (nodes are organized into layers that accept input, transform it, and pass it along; traditional neural networks may have an <i>input</i> layer, an <i>output</i> layer, and one or more <i>hidden</i> layers in between)
Depth	The number of layers in a neural network
Deep learning	ML using neural networks with more than one hidden layer
Feedforward neural network	A type of neural network in which values are passed from the input layer, through hidden layers, to the output layer, and are subject to transformation at each layer
Gradient descent	A common optimization method used to train neural networks. At each step, the parameters of the model are adjusted a small amount in the direction of greatest decrease in the error function. This is akin to charting a trajectory down a mountain by always moving in the direction of the steepest descent. Note that the resulting path may not be shortest and may not converge to the global minimum.
Backpropagation	A widespread method that uses the chain-rule from Calculus to calculate the gradient of a neural network's error function with respect to its weight parameters. Often used in conjunction with gradient descent to train neural networks
Convolutional neural network	A type of feedforward neural network that was originally developed for processing images that uses specially learned convolutions or "filters" to exploit certain local structures in data
Recurrent neural network	A family of neural networks structured to process inputs of arbitrary length
Autoencoder	An unsupervised method primarily used for noise filtering
Random forest	A different kind of supervised ML technique that uses an ensemble of decision trees to classify an input. Random forests have several advantages over neural networks, including being relatively human-comprehensible and requiring relatively little tuning, but usually are only used as classifiers on relatively low-dimensional data.
Bayesian learning	A framework for ML in which one starts with an initial guess, or "prior" for the probability distribution in question, and then the estimate for that distribution is updated according to Bayes' theorem as new observations are made.
Hidden Markov Model	A statistical model in which the system being modeled is assumed to be a Markov process (i.e., the probability of a future state depends only on the immediately prior state), but some or all of these current and future states may be unobserved ("hidden")

Current Status

AI and ML are broad, diverse fields encompassing a variety of techniques, many of which may be applicable to basic analytic tasks. Today’s ML algorithms are very powerful “function approximators,” or systems that are trained to map input data to output data and then predict what outputs would be derived from different (untested) input data. There are two main flavors of ML today: supervised and unsupervised. The former must learn from labeled training data, while the latter seeks to learn from training data without labels. A third flavor, reinforcement learning, lies somewhere between the two: Instead of learning from explicit labels, the algorithm attempts to maximize some reward signal by interacting with its environment. We discuss reinforcement learning in more detail later in this chapter.

ML has achieved remarkable empirical results in solving certain types of problems, including machine vision, translation, and speech recognition. However, despite demonstrated successes, difficulties remain. Many ML algorithms require large amounts of training data, which often must be labeled by humans or acquired in some other costly fashion. Often, cleaning or formatting real-world data to create a usable training set makes up a significant portion of the work of building a practical ML application. For example, Project Maven works with existing AI/ML algorithms but requires a training set of *millions* of hand-labeled images.¹³⁰

Another challenge that persists even with state-of-the-art methods is that of “overfitting” the training data—meaning that the algorithm is exquisitely tuned to the data on which it was trained but generalizes poorly to real-world data beyond it. Although we ideally would have ML techniques that could be applied to arbitrary problems without additional tinkering, this goal remains a long way off, if it is possible at all. Most current ML methods require many different parameters to be finely tuned, and finding the right combination of parameters to get the algorithm to work can be forbiddingly difficult.

Future Development

Compared with other aspects of computing, such as the growth in the density of integrated circuits, for which Moore’s Law proved accurate for 50 years,¹³¹ the future development of AI has been notoriously difficult to predict. Many early AI researchers mistakenly thought that “human-level” AI would be created well before the start of the 21st century. Then the pendulum

¹³⁰ The ML algorithm itself has been described as just “75 lines of Python code” on top of Google’s standard TensorFlow package, but it is estimated that the project will need to generate at least 100,000 individually labeled images for *each* of the 38 different object types it hopes to recognize (Cheryl Pellerin, “Project Maven to Deploy Computer Algorithms to War Zone by Year’s End,” press release, Washington, D.C.: U.S. Department of Defense, July 21, 2017; Lynette M. Role, “New Artificial Intelligence Technology Assists Air Commandos with Decision-Making,” press release, Hurlburt Field, Fla.: Air Force Special Operations Command Public Affairs, September 13, 2017).

¹³¹ Max Roser and Hannah Ritchie, “Technological Progress,” Our World in Data webpage, 2013.

swung the other way and algorithms, such as deep learning and reinforcement learning, were dismissed as failures by much of the research community in the 1990s and early 2000s. As recently as ten years ago, the imminent conquest of AI and ML by deep learning was far from obvious, even to most leading researchers.

Still, surveys of experts vary enormously in this area. For example, a 2016 survey on when AI would be able to succeed at certain tasks reported estimates that varied wildly between ten and 50 years for complex tasks, such as writing a novel. Even for seemingly simple tasks, such as folding laundry, expert estimates varied by more than ten years.¹³² It is therefore extremely difficult to predict the future state of the field, and it is not given that future breakthroughs will build on currently dominant techniques. As a consequence, it is difficult to predict when AI/ML techniques will be able to perform the more complex analytic tasks, particularly those that require a degree of understanding or common sense, as opposed to mere statistical pattern-matching.

Current techniques either rely on statistical inference or can reason using a provided model but do not “understand” in the sense needed for many tasks. For example, it would be desirable to have scene-analysis systems that are intelligent enough to anticipate *why* a particular confluence of possibly mundane observations was significant and then alert human analysts. Current techniques, however, would generalize poorly for such purposes unless the training set anticipated it. Problems requiring this sort of “human-like” or “general” intelligence are informally referred to as “AI-complete” among researchers.¹³³

In contrast to the hazy trajectory of AI and ML research, it *is* possible to anticipate commercialization of existing technologies with greater confidence. Major technology firms are aggressively seeking ways to scale up recent research findings to address larger data sets. These efforts include both developing larger training sets, such as labeled images and multilingual texts, and developing custom hardware for training and executing ML models. Improvements in software engineering can provide similar speedups. As these advances are commercialized, current ML technology will be able to address more and larger real-world tasks even without additional breakthroughs in ML algorithms or architectures. Although the requisite information

¹³² Estimates by AI researchers for what they evidently deemed to be the most complex task, automating AI research itself, varied by more than a century (see Katja Grace, John Salvatier, Allan Dafoe, Baobao Zhang, and Owain Evans, “When Will AI Exceed Human Performance? Evidence from AI Experts,” *Journal of Artificial Intelligence Research*, Vol. 62, 2018; see discussion in AI Impacts, “2016 Expert Survey on Progress in AI: Narrow Tasks,” webpage, undated).

¹³³ The term *AI complete* was coined in analogy to “NP complete” from mathematics. An NP complete computing problem is at least as difficult as any other computing problem for which an answer can be checked in polynomial time. If there were a method to solve any NP complete problem in polynomial time (although no such method is yet known), that method could be used to solve any computing problem in polynomial time, as long as the problem permits answers to be checked in polynomial time. By analogy, an AI complete problem is difficult enough that, if AI technology could solve it, then that same technology could be adapted to do anything else that human intelligence can do (Roman Yampolskiy, “Turing Test as a Defining Feature of AI-Completeness,” in Xin-She Yang, ed., *Artificial Intelligence, Evolutionary Computing and Metaheuristics*, Berlin: Springer-Verlag, 2012).

to analyze the resultant performance boost from hardware is often proprietary, it should be calculable in a way that improvements from additional data or better pretrained models are not. The extent to which this process will enhance individual applications is necessarily specific to that task.

We now take a closer look at several commonly used AI/ML methods and algorithms.

Artificial Intelligence/Machine Learning Methods


Neural Networks

Artificial neural networks are one kind of ML and are loosely inspired by the way that neurons are linked together in biological brains. These networks map inputs to outputs and are trained to build maps through various learning processes. The inputs and outputs to a neural network usually take different forms. For instance, a neural network that is trained to recognize bananas might accept an *image* as its input and might return as its output a *number* representing its confidence that the image indeed depicts a banana. During the training phase, a neural network would initially output random numbers but would learn gradually to distinguish the images that contain bananas, reducing the error rate after many tries. If trained well and supplied with sufficient input data, neural networks can be effective. They can even be predictive. For instance, neural networks used in today's facial-recognition systems can estimate what a known face would look like from a previously unseen angle.

Neural networks are structured in layers, as shown in Figure 4.1. In most structures, each layer accepts an input from the previous layer, transforms it, and passes its output to the next layer. Networks that follow this structure are also called *feedforward networks*. The output from each node is usually a weighted average of the inputs, normalized to within some standard range (such as between 0 and 1). The process of training a neural network is essentially the process of determining these weights.

Most neural networks will have at least an input and an output layer. The *depth* of a neural network refers to the number of “hidden” layers between these two layers. A neural network with one or more hidden layers is said to allow for *deep learning*. The way in which these layers are wired together is one of the main distinctions among different types of neural networks. Neural networks may contain many layers with a large number of nodes in each.

Figure 4.1. Layered Structure of Neural Networks



Despite the current excitement around deep learning, neural networks are often difficult to implement in practice. Applying neural networks to real-world problems typically requires considerable trial and error on the part of the developer. Neural networks are also comparatively resource-intensive, in terms of both computational resources and training data, and these requirements can increase rapidly with the depth of the network.

Furthermore, neural networks are “black boxes” in the sense that their behavior is difficult to interpret. For example, if an algorithm misclassifies something as a banana, it is often very difficult to find out why. Thus, it is preferable to use simpler methods where feasible. Older ML techniques, such as *random forests*, need less tinkering, are easier for humans to understand, and give good results in many cases.¹³⁴ But for many desired applications, such as machine vision, deep learning is the best (and perhaps only) currently viable technique.

Decades of research have resulted in a wide variety of different types of neural networks, although only a subset of these are in widespread commercial use. In current practice, it is unusual for a single type of neural network to be employed in isolation in an end-user application. Instead, they are combined with other kinds of neural networks and/or nonneural network components into larger architectures.

We review different types of neural networks and related methods in the next sections.

¹³⁴ *Random forests* are an ensemble-learning technique using decision trees. Because of the distribution of models learned by the component decision trees, their errors tend to cancel each other out, resulting in an overall model that generalizes well to unseen cases. Although they are not suitable for such tasks as machine vision, random forests are relatively easy to use compared with neural networks. They have few parameters to tune and often “just work” when applied to an appropriate classification task.

Fully Connected Neural Networks

A neural network is considered fully connected if each node connects to every node in the previous layer and every node in the next layer. (The network shown in Figure 4.1 is fully connected.) The classic fully connected feedforward neural network was developed in the 1980s as an enhancement to the much-older perceptron.¹³⁵ *Perceptrons* are two-layer neural networks that were originally implemented as custom hardware. In 1969, however, Minsky and Seymour Papert showed that the perceptron was theoretically incapable of learning many simple functions, effectively squelching interest in them until the mid-1980s.¹³⁶

At that time, researchers began to add hidden layers. Computational techniques, such as backpropagation and stochastic gradient descent, were breakthroughs in training these neural networks. Stochastic gradient descent uses a “hot-or-cold” response to measure how close an output is to the correct answer for a given example in the training set; *backpropagation* is a straightforward technique of working backward through the layers to adjust their weighting parameters so the next guess would be closer to the mark.¹³⁷ Researchers proved that these deeper neural networks can, in theory, learn to approximate *any* continuous function given sufficient training examples.¹³⁸ These techniques are the foundation of most neural networks today, and, even with the advent of deep learning techniques, fully connected neural networks are currently at the core of many technologies.

However, because they are fully connected, their computational requirements grow substantially with the size of each layer. Some tasks grow exponentially.¹³⁹ This is especially challenging for images, which can be very large. Furthermore, their inputs and outputs are fixed in size, which, consequently, makes such networks difficult to adapt to text, audio, or video formats where records vary greatly in length. To address some of these limitations, researchers developed alternatives, such as convolutional neural networks (CNNs) and recurrent neural networks (RNNs), which are both discussed in the next sections.

Given their widespread employment in industry and research applications, such as those using Google’s TensorFlow, fully connected neural networks are mature. The applications built using these techniques, however, need not be.

¹³⁵ David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams, “Learning Representations by Back-Propagating Errors,” *Nature*, Vol. 323, October 9, 1986.

¹³⁶ Marvin Minsky and Seymour Papert, *Perceptrons. An Introduction to Computational Geometry*, Cambridge, Mass.: MIT Press, 1969.

¹³⁷ Of course, this method only works for problems where one can define being “close” to the right answer. Doing this for image recognition, for example, may require some thought when one is literally comparing images of apples to oranges.

¹³⁸ See, for example, Kurt Hornik, “Approximation Capabilities of Multilayer Feedforward Networks,” *Neural Networks*, Vol. 4, No. 2, 1991.

¹³⁹ Many heuristics and approximate methods are used, however, to keep these requirements down.

Convolutional Neural Networks

CNNs are a derivative of classic feedforward neural networks that are designed to take advantage of local structure in data. These are not fully connected: Each node accepts input only from a subset of the previous layer. CNNs were originally developed to process images but have also been employed successfully for such tasks as machine-language translation and signal processing. As an alternative to previous machine-vision approaches that employed hand-engineered visual features, CNNs can learn features automatically in an end-to-end fashion. They do this by employing layers that consist of volumes of neurons to represent such concepts as depth, height, and width.

The power of convolutional networks lies in their ability to learn the convolutions from data using backpropagation, just like a fully connected network. Once trained, the CNN produces a “learned representation” that converts an input, such as an image, into a fixed-size vector. In typical applications, these learned representations are then used as input to a fully connected neural network.¹⁴⁰

CNNs have rapidly supplanted earlier methods in most machine-vision tasks.¹⁴¹ CNNs now often feature hundreds or even thousands of layers. A logical counterpart of the CNN is the deconvolutional neural network, which works on the same principles, only in reverse. For instance, a deconvolutional neural network can take a fixed-length representation vector and use it to produce an image. This technique is used to make deepfakes, which will be discussed later in this section.

Training a high-quality CNN with a substantial training set can be a forbidding endeavor, but fortunately, a CNN can often be repurposed for a new application by simply retraining its last few layers. In machine-vision systems, the earlier layers of the CNN typically represent low-level image features, such as edges and corners, that generalize well between use cases. High-level features corresponding to different kinds of objects occur in late convolutional layers.

Like all forms of deep learning, CNNs are very data hungry and can require large amounts of central processing unit time to train. Furthermore, CNNs are a supervised ML technique that require labeled training data. Moreover, they are still “black box” techniques with poor interpretability and do not allow for explicit transfer learning (i.e., the ability to transfer what has been learned from one data set to another easily). They can also be brittle and give incorrect or nonsensical outputs when presented with images that are dissimilar from those in their original training set.

Because CNNs are employed for almost all state-of-the-art machine-vision tasks, including object and facial recognition, we consider them to be mature.

¹⁴⁰ Yoshua Bengio, Aaron Courville, and Pascal Vincent, “Representation Learning: A Review and New Perspectives,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 35, No. 8, August 2013.

¹⁴¹ Krizhevsky, Sutskever, and Hinton, 2012.

Recurrent Neural Networks

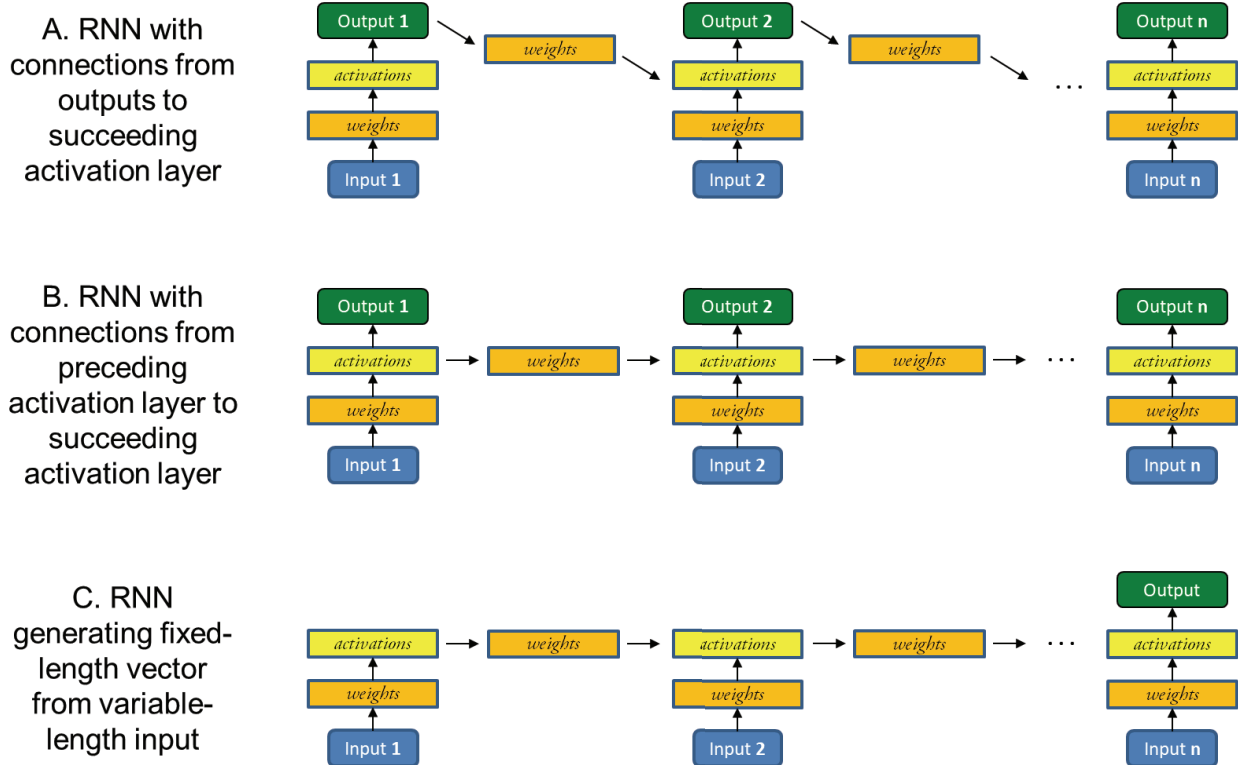
RNNs are a family of neural networks that are structured differently so they can process arbitrary-length sequences, such as audio, video, and radar data. Because input data often vary in size and length, RNNs are an essential part of the ML expert’s tool kit.

Typical RNNs share characteristic features with feedforward networks, including the use of weights, activation functions, hidden layers, and backpropagation. The difference lies in the existence of additional connections that allow the propagation of information forward in time and the use of special activation functions that try to retain “memories” of important earlier inputs until they are needed. The presence of this “cycle” allows them to, in effect, execute a learned computer program with certain fixed parameters. In theory, RNNs can learn to mimic almost any computable function.¹⁴² RNNs are typically trained using backpropagation through time, which is like ordinary backpropagation, except that losses are propagated “back in time” through the unrolled network.

The recurrent connections in RNNs can take various forms, as illustrated in the unrolled diagrams in Figure 4.2. A common arrangement is to have the activations function consider not just the weighted inputs but also the weighted outputs from the previous time step (A). A more powerful, albeit harder to train, arrangement has recurrent connections between the activations layer of each time step (B). Using the previous activations rather than the previous output as an input to the activations layer can preserve information that would otherwise be lost. When desired, RNNs can translate a sequence to a fixed-length vector representation (C). This fixed-length vector can be used as an input to CNNs or other feedforward neural networks—for example, such an RNN could translate an English-language description into a vector that a deconvolutional network or generative adaptive network (discussed on the next page) would then use to generate an image.

¹⁴² Hava T. Siegelmann and Eduardo D. Sontag, “On the Computational Power of Neural Nets,” in *Proceedings of the Fifth Annual Workshop on Computational Learning Theory*, Pittsburgh, Pa., 1992.

Figure 4.2. Recurrent Neural Networks



SOURCE: RAND analysis.

There are many kinds of RNNs, but a common component is long short-term memory (LSTMs) units, which were introduced by Jürgen Schmidhuber in 1997. LSTMs were invented to avoid the “catastrophic forgetting” problem that bedeviled early RNNs—the relevant earlier information was diluted as it was propagated forward in time, so it tended to be lost by the time it was needed by the subsequent nodes. The LSTM employs a memory cell that can store data to control how the networks unroll, enabling it to look forward or backward in the input sequence, depending on its previous inputs. The LSTM can be used in combination with other kinds of units within an RNN, and RNNs using it are commonly employed in systems employing other techniques, such as CNNs.

RNNs are employed in state-of-the-art machine-language translation and speech-recognition systems. They are also used for text summarization and scene description (typically in conjunction with CNNs). RNNs can also be used as an attention mechanism to seek certain kinds of features in imagery, “steering” the system to where objects of interest are likely to be located. For example, the current version of the open-source Tesseract optical character-recognition library exploits LSTMs to detect where lines of text are located in the page.¹⁴³ This same

¹⁴³ GitHub, “tesseract-ocr,” webpage, undated.

principle can be exploited for spatial and temporal features in other kinds of data, such as audio and video.

Although they require less hand-engineering than previous techniques for applications, such as machine-language translation, RNNs are still challenging to design and train. Moreover, the types of RNNs needed would be supervised systems requiring extensive labeled training data, such as translation systems requiring bilingual texts for training. Although systems using RNNs can produce eerily convincing results, they are still “black box” techniques and can show brittleness when provided with unfamiliar inputs. This makes them difficult to validate and debug.

Recurrent neural networks, such as LSTMs, are in widespread commercial use in industry.

Autoencoders

Autoencoders are neural networks that learn an effective, reduced-size internal representation of the input data that can be used to reconstruct it as faithfully as possible. One prominent example is the denoising autoencoder (DAE). As their name implies, DAEs can be employed as noise filters; when trained effectively, they learn to focus on preserving the “essential” features of the input. Deep autoencoders can even learn to make educated guesses about the signal hidden behind noise in the original.¹⁴⁴ DAEs therefore have considerable utility for data processing and compression tasks. They could also be deployed on sensor platforms to compress data in an intelligent, adaptive way to make the most of available transmission bandwidth.

As an unsupervised learning technique (in the sense that DAEs learn to reconstruct their training samples), they require less labeled training data. The network is trained using samples with added noise to teach it to distinguish the “true” signal. Despite the simplicity of this scheme, their effectiveness is still highly dependent on the choice of network architecture and training data. Like other kinds of neural networks, they are liable to be brittle in the face of inputs too far outside the training set. If DAEs are not trained on a data set that is similar to the real-world use case, they may generate an internal representation that is ill-equipped to capture unusual features in the environment. As a consequence, under some circumstances, they may filter out atypical and interesting features that are mischaracterized as “noise.” DAEs are widespread in commercial use in industry.

Deep Reinforcement Learning

Reinforcement learning aims to create agents that act in an environment to maximize some reward or utility function. This is distinct from supervised learning in that “correct” answers are never provided during the training process, only punishments and rewards. Although reinforcement learning predates deep learning, over the past few years, the combination of the

¹⁴⁴ Because of this capability, they can be used to generate “deepfakes,” albeit of lower quality than Generative Adversarial Networks—see the section on adversarial examples later in this chapter.

two has led to dramatic progress on certain tasks, particularly game playing. The most dramatic examples of this are DeepMind's Alpha* agents (AlphaGo, AlphaGo Zero, and AlphaZero), along with its agents that learn to play Atari 2600 games from raw pixels.¹⁴⁵ The application of reinforcement learning to help tune or design deep learning systems for other tasks is also an active area of research.¹⁴⁶

The idea of reinforcement learning is quite general, but the field's most spectacular successes came when it was combined with deep neural networks, leading to so-called deep-reinforcement learning. Although deep-reinforcement learning is currently the state-of-the-art approach for game playing, the typical approach has some limitations that have hindered attempts to apply it to other tasks. It is extraordinarily sample-inefficient, requiring enormous amounts of experience merely to learn the nature of a task, much less to achieve superhuman performance. In some domains, such as games, this obstacle can be overcome by simulating play at high speed and in parallel. But for many domains where real-world input is needed, such as with robotics, low sample efficiency is a serious challenge. This is an area of intense research.

Deep-reinforcement learning is currently technically challenging, even for talented AI researchers, and its relative technological immaturity may be a greater obstacle to adoption by the Air Force than for commercial users because of a low tolerance for error in certain applications. In a few years, however, research might advance to the point that it can contribute to the various AF DCGS processes.¹⁴⁷

The current technological maturity of deep-reinforcement learning depends on the type of system being considered and its application. Some game-playing systems, such as AlphaZero, are state-of-the-art but not yet commercialized. Google has announced that it is using DeepMind reinforcement learning to manage energy usage in its server farms. A particular challenge to assessing the technological maturity of reinforcement learning is that current excitement over the technology has resulted in a tendency to label systems that are really variants of supervised learning as "reinforcement learning."

Supporting Technologies

Applying AI/ML to AF DCGS operations requires more than just algorithms: One must also manage the data to feed them and have the hardware sufficient to run them. This section offers a brief overview of some of these necessary technological enablers that support AI/ML. These

¹⁴⁵ Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dhharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis, "Human-Level Control Through Deep Reinforcement Learning," *Nature*, Vol. 518, No. 7540, February 25, 2015, p. 529.

¹⁴⁶ Barret Zoph and Quoc V. Le, "Neural Architecture Search with Reinforcement Learning," talk presented at the 2017 International Conference on Learning Representations, Toulon, France, 2017.

¹⁴⁷ This technique may also be applicable to the operation of ISR platforms (i.e., ability to conduct evasive maneuvers while out of contact with remote operators).

technologies, of course, also have significant benefits in terms of efficiency and effectiveness well beyond their ability to facilitate subsequent AI/ML introduction.

Data Management and Storage

Because AI/ML algorithms are data-intensive, their practical use depends on effective management of training data. Databases are needed to store intelligence data and findings in a convenient form so that analysts can query data and the databases can serve as training samples for ML algorithms.¹⁴⁸ Mature database solutions exist for a wide variety of use cases. Large-scale databases have been in widespread use for decades. Although relational databases founded on Structured Query Language (SQL) are the most common (e.g., MySQL and Microsoft Access), alternatives, such as graph databases and document databases, are widely employed for industrial purposes. Next-generation solutions are likely to be expensive and disruptive and need to be acquired with attention paid to cost-effectiveness, performance, and security. Enabling multilevel security and the appropriate compartmentalization of classified information is also a vital consideration. Not only does secure information need to be strongly encrypted, it also needs to be isolated on appropriately secure systems for policy compliance.

Similarly, the Air Force needs a good solution to store the enormous amount of data it collects even before these data are sorted into a database. Although much ISR data today is stored in national databases, the AF DCGS will likely require a way to store its own raw data from ISR platforms for the purposes of training ML algorithms. Historically, tape has been used for large-scale storage solutions because it has been cheaper than solid-state drives (SSDs) or disks, but tapes must be warehoused and can be accessed only by physically retrieving the stored media. Most contemporary commercial applications have now moved to SSDs for improved performance, and, as prices continue to fall, the cost differential soon may be negligible.¹⁴⁹ As with databases, the ultimate large-scale data-storage solution needs to be cost-effective and secure (whether via encryption, physical security, or both).

Cloud storage and computing solutions are increasingly ubiquitous in the commercial sector. The AF DCGS will need to decide how to transition to large-scale cloud computing. One possibility is for DoD to develop in-house cloud computing capabilities, but this option may prove very expensive relative to its performance. Another option will be to work with commercial solutions that are tailored to its needs.

Custom Hardware

Even once the AF DCGS has determined the nature of its next-generation computer-networking and data-storage solution (cloud or otherwise), there remains the issue of what type

¹⁴⁸ To date, this has been implemented with varying levels of success.

¹⁴⁹ See, for example, Andy Patrizio, “SSDs Get Bigger, While Prices Get Smaller,” *Network World*, May 22, 2018.

of chips will be harnessed, especially for ML. There has been a revolution in the past few years in the development of special computer chips designed for ML applications. We review two examples in the next two subsections.

Neuromorphic Computing

Neuromorphic computing aims to create computer hardware that works “like the brain,” or, in many cases, to train or evaluate deep neural networks efficiently. This is often essentially a marketing term. However, a major objective of most of these devices is to operate in a more energy-efficient way than equivalent central processing units and graphics processing units (GPUs). Neuromorphic components might eventually be used by the Air Force for training and executing neural networks, or they could be integrated into sensor platforms as a way of conserving bandwidth. Presently, however, these technologies are too diverse to make any general assessment about their maturity for these purposes.

Tensor Processing Units

Tensor processing units (TPUs) are custom computer hardware designed by Google for matrix operations, which form the basis of modern neural network implementations.¹⁵⁰ These proprietary devices are fairly power hungry and are not currently sold to end users but rather are leased via the cloud. Such devices as TPUs make the training and evaluation of modern neural networks much faster than other kinds of hardware, including high-end GPUs. The increasing number of firms seeking to commercialize similar technologies offers hope that equivalent capabilities will soon be available, and the AF DCGS should have several options to choose from in the future. TPUs are already in use at Google’s data centers.

Vulnerabilities

As just discussed, AI/ML methods have many potential applications for Air Force intelligence. However, we should also be mindful that new technologies might come with new vulnerabilities. There is an increasing awareness that many ML algorithms can be “tricked” by what are commonly referred to as *adversarial attacks*.¹⁵¹


There are different types of adversarial attacks. A common one is an input that is specifically designed to confuse an AI/ML model and cause it to produce incorrect results. Recent research has shown that current instantiations of machine-vision systems, such as object classifiers, are highly vulnerable to adversarial attacks. These are especially important in cases in which the

¹⁵⁰ Cade Metz, “Google Makes Its Special A.I. Chips Available to Others,” *New York Times*, February 12, 2018.

¹⁵¹ RAND researchers are working on a growing body of work on this subject. For example, see Gavin S. Harnett, Andrew J. Lohn, and Alexander P. Sedlack, “Adversarial Examples for Cost-Sensitive Classifiers,” working paper presented at the 33rd Conference on Neural Information Processing System (NeurIPS 2019), Vancouver, Canada, December 13, 2019.

human eye cannot discern adversarial attacks, such as in the famous example shown in Figure 4.3.¹⁵² In this figure, we see that a small amount of what appears to be random noise has been added to the picture. This extremely subtle distortion, invisible to the human eye, can fool the AI system into misclassifying the object: The algorithm has mistaken a panda for a gibbon (in other cases, a turtle has been mistaken for a gun¹⁵³). This poses a major concern because adversaries can be expected to attempt such techniques to counter U.S. intelligence applications of machine vision.

Figure 4.3. Tricking Machine Vision



SOURCE: Goodfellow, Ian, Jonathon Shlens, and Christian Szegedy, "Explaining and Harnessing Adversarial Examples," paper presented at the International Conference on Learning Representations 2015, May 7–9, 2015, p. 3.

Generally speaking, machine vision systems using CNNs are known to be vulnerable to adversarial examples. RNNs are also vulnerable to adversarial inputs, although work to date in this area is less mature. Although the most-discussed examples of such attacks were designed to confuse image classifiers into giving specific incorrect results, it is also possible to design inputs that are designed to lower the accuracy of the classifier without any particular intended outcome.¹⁵⁴ Researchers have demonstrated empirically that adversarial examples have some counterintuitive qualities. Designing them requires relatively little access to the classifier, and they can work on multiple classifiers, particularly those that share architecture and/or training data. The common use of CNNs pretrained on public data sets increases these vulnerabilities, but it is not their sole cause. The recent demonstration of adversarial examples that confuse time-

¹⁵² Goodfellow, Shlens, and Szegedy, 2015.

¹⁵³ Abigail Beall, "Visual Trick Fools AI into Thinking a Turtle Is Really a Rifle," *New Scientist*, December 3, 2017.

¹⁵⁴ Justin Gilmer, Ryan P. Adams, Ian Goodfellow, David Andersen, and George E. Dahl, "Motivating the Rules of the Game for Adversarial Example Research," *arXiv preprint arXiv:1807.06732*, July 20, 2018.

limited humans as well as machine-vision systems suggests that defense against adversarial inputs may be extremely difficult even for powerful future AI systems.¹⁵⁵

In a different type of adversarial attack, adversaries could defeat hidden Markov models, such as those used in speech recognition, by intentionally violating the assumptions of the underlying pronunciation and language models. For instance, speaking pig Latin in a peculiar accent could cause an otherwise well-performing model to output nonsense. Indeed, researchers have demonstrated adversarial attacks on speech-recognition systems.¹⁵⁶ This even extends to more-advanced methods, such as deep-reinforcement learning.¹⁵⁷

Despite the amusing examples that we use to illustrate adversarial attacks, this represents a real problem for AI/ML-based analysis. Removing the human analyst from the loop entirely would leave the system exposed to adversarial attacks of this kind. The Air Force will need to consider active defenses to adversarial attacks in its ML implementations from the outset of development and going forward. Although there is disagreement in the field about how well researchers will be able to harden their algorithms against these attacks, anyone using ML algorithms should be aware of the variety and effectiveness of existing adversarial attacks and keep tabs on how they evolve.

¹⁵⁵ Gamaleldin Elsayed, Shreya Shankar, Brian Cheung, Nicolas Papernot, Alex Kurakin, Ian Goodfellow, and Jascha Sohl-Dickstein, “Adversarial Examples That Fool Both Computer Vision and Time-Limited Humans,” paper presented at the 32nd Conference on Neural Information Processing Systems, Montréal, Canada, December 4, 2018.

¹⁵⁶ Nicholas Carlini and David Wagner, “Audio Adversarial Examples: Targeted Attacks on Speech-to-Text,” presented at the 39th IEEE Symposium on Security and Privacy, San Francisco, Calif., May 24, 2018.

¹⁵⁷ Sandy Huang, Nicolas Papernot, Ian Goodfellow, Yan Duan, and Pieter Abbeel, “Adversarial Attacks on Neural Network Policies,” workshop paper presented at the Fifth International Conference on Learning Representations, Toulon, France, April 26, 2017.

5. Improving GEOINT Analysis: Additional Detail

In Volume 1, we recommended short-term (using today's technology) and longer-term (using future technologies as they mature) applications for AI/ML in the areas of GEOINT, SIGINT, OSINT, multi-INT, networking and hardware, and mission management.¹⁵⁸ This chapter presents additional details on the GEOINT recommendations and how they would fit into the generic data flow map first presented in Chapter 2. As in Volume 1, we divide the discussion into improvements using today's technology and those using future technologies as they mature.

Making the Most of Today's Technology to Improve GEOINT Analysis

Table 5.1 reproduces the short-term GEOINT recommendations that were outlined in Volume 1, the analysis roles that would be affected, and the objectives they would help address.

Table 5.1. GEOINT Recommendations: Making the Most of Today's Technology

Summary of Recommendation	Analysis Roles Affected	Objectives Addressed
Create a geospatial intelligence analysis and reporting tool (GEOART) to semiautomate product generation and mission reporting.	Reporter	Efficiency Human capital
Create an improved formatter to assist with threat warning.	Reporter	Efficiency
Create a linker tool to tie information used to confirm the exploitation back to the source.	Exploiter Investigator	Effectiveness
Adopt geographic information system (GIS) into the MTI workflow.	Exploiter Investigator	Effectiveness Agility
Bring in programmers to write Python scripts to automate analysis processes within GIS.	Exploiter Reporter	Efficiency
Assess risks and benefits of adopting the industry standard in video-editing tools for WAMI.	Exploiter Investigator	Effectiveness Agility

The primary goal of adding these tools is to significantly speed up GEOINT analysis processes by preventing duplication of effort and standardizing the workflow. Today's processes include significant cutting and pasting, manual entry of metadata from one system to another, and manual preparation of many products. For high-altitude imagery, most of the analyst's time is spent formatting, not analyzing. So, for example, because object-recognition tasks take little time, automating them with AI/ML will not add efficiency to AF DCGS operations unless

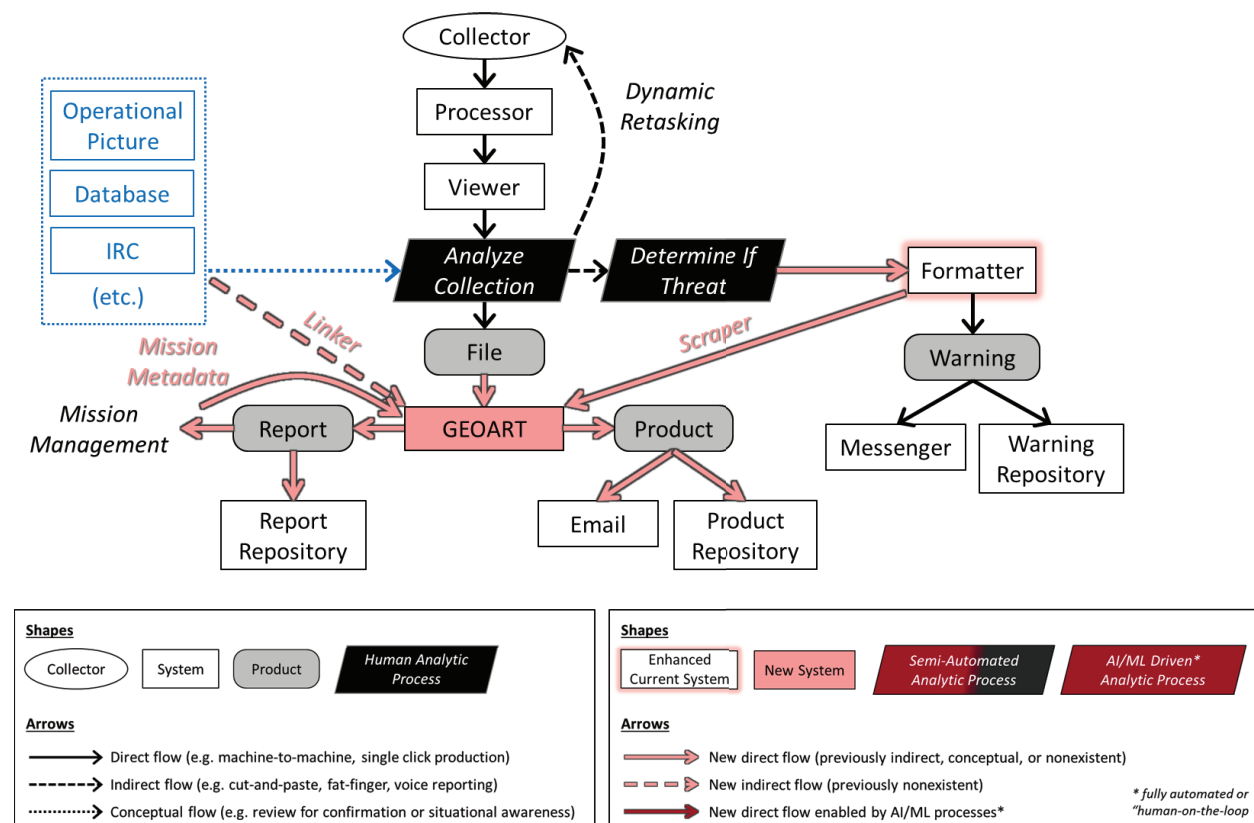
¹⁵⁸ Menthe et al., 2021.

reporting tasks are also automated; automation might instead exacerbate existing bottlenecks in the reporter role. Such automation could, however, allow for archival imagery to be tagged with metadata, a task for which there are few resources available today. In addition to speeding up existing processes, many of the technologies listed in the next sections would prepare the way for the more-advanced future of AI/ML technologies that we will discuss later in this report.

GEOINT Analysis and Reporting Tool

A primary recommendation that can be implemented today is to create a GEOART to generate reports for all high-altitude imagery, FMV, and MTI missions. Details of how a GEOART would plug into the existing architecture are discussed in the restricted volume.¹⁵⁹ Figure 5.1 shows how GEOART (bottom center) and the other nterm recommendations discussed later would transform the generic data flow map introduced in Chapter 2.

Figure 5.1. Generic Data Flow Map—Potential Improvements with Today’s Technology



As it is presently done, the exploiters would select images or screenshots and save them to a folder. The GEOART would then automatically format PowerPoint slides with appropriate

¹⁵⁹ Menthe et al., forthcoming.

templates for whatever products are required for the mission, such as a vehicle follow report. At this stage of development, products would likely still require some additional hand annotation (depending on the INT), but most of the required data and formatting could be added automatically, which would save considerable time for human analysts. This change in the workflow would prepare for future developments as well. As AI/ML methods permit more-automated tagging of information, this information would be made available to the GEOART, which would require less human intervention to complete reports. Eventually, we envision that this would become a “human-on-the-loop” process that could accept human intervention but would not require it. In addition to formatting the product, the GEOART would also format a mission report (e.g., target log) using these images, the goal being to ensure that the analyst does not need to enter information twice.

Formatter for Threat Warning

The next recommendation, shown on the right side of Figure 5.1, is to create an improved formatter to assist with threat warning. Instead of hand-jamming information into a web-enabled interface, the analyst would need only to push a button. Relevant metadata concerning the warning—source, method, parameters—would be properly formatted and pushed out over the warning systems. This would be particularly helpful for real-time intelligence, where warnings must be repeatedly updated by hand today. Some analysts we spoke to expressed concern that, without such an improvement, it would be difficult to scale up certain warning processes to handle what might be required in a large-scale conflict.

We expect that the main difficulty in creating such a formatter would be in crossing classification barriers and dealing with outdated hardware systems. Indeed, were there no such barriers, airmen likely would have created such a formatter already.¹⁶⁰ Open-architecture DCGS might alleviate some of these barriers, but it certainly will not eliminate all of them. Again, one purpose of making this change is to prepare the way for a semiautomated threat-detection process that would send warning messages automatically. Creating an improved formatter today would not be a wasted effort but would help address “hidden” implementation barriers that could block future AI/ML upgrades if not solved.

Linker

The next recommendation is to create a linker, shown as a dotted arrow on the center-left side of Figure 5.1, that would prepare future back-tracing of exploitation to collections. Presently, a lot of reference material is reviewed to confirm exploitation, but records of what

¹⁶⁰ We heard of a few efforts to automate different kinds of warning messages, in fact. Many of these eventually died when the enthusiastic airman who invented them left the AF DCGS, but others may remain viable. The weapons and tactics (WEPTAC) conference we discuss in Chapter 7 would be a good forum for soliciting information on this. The technical problems may already be largely solved.

information was reviewed are rarely kept because doing so would add extra steps to the analysis process and is not currently required. If this process can be simplified, however, such as by storing a uniform resource locator (URL) or other reference with a single push of a button, then these references could be integrated into the final report—and would pave the way for more-automated linking in future development cycles.

Automatically identifying all relevant contextual data from various systems would need to wait for AI/ML tools. Identifying geolocated data from the area of interest would be simpler but would still be complex. In the future, when these processes are semi- or fully automated, the data pulled by an AI/ML tool from reference databases would follow the same paths and provide new traceability to show how collections link together. Thus, solving the short-term problem of linking information between separate systems would not be wasted effort; indeed, it is necessary to set the conditions for successful integration of AI/ML-enabled workflows.

Geographic Information System for MTI

Initially, the GEOART just described would be more directly applicable to high-altitude imagery and FMV. Ultimately, we would like GEOART to extend to MTI as well, mostly so that those missions and their products can be reflected in the Unified Collections Operations Reporting Network (UNICORN) architecture like the rest of GEOINT. However, because MTI products are not as standardized (they depend more on theater needs), this aim can wait for a later revision of UNICORN. Initially, for MTI, our primary recommendation is to proceed with integrating the GIS into the process. GIS is a program that is designed to store, analyze, and display georeferenced data, typically by combining information in layers that can be manipulated separately. We believe that such a tool could replace *all* the manipulations currently done with various custom tools and Excel (different tools are used at different DGS sites).

The phenomenology of MTI, which involves visualizing a forest of geolocated point returns and movement vectors (“dot-ology”), is uniquely suited to GIS work. Furthermore, modern GIS software natively accepts Python scripting to automate repetitive processes. Writing these scripts should not be difficult. One airman we spoke to was learning to code on his own time so he could write these scripts for GIS systems in the future. When asked whether he thought a professional programmer could write these scripts within 30 days, he laughed and said, “Yes, but what would he do with the other 25 days?”

Integrating a complex system, such as the GIS, into any part of the AF DCGS ops floor would require additional training, however. Some consideration would need to be given to when and where this training would occur. As discussed in Chapter 7, special certification programs for interested airmen would enable them to be qualified not just to use the system but also to teach it to others.

It is true that such a system as GIS is arguably more powerful than the AF DCGS needs today. But the GIS environment is the standard for the IC, notably the NGA.¹⁶¹ The 11th Special Operations Intelligence Squadron and other organizations are already moving in this direction. MTI can be the spearhead that helps introduce GIS more fully into AF DCGS operations, particularly at the DART, where the kind of geolocation and data fusion it does is also well-suited to GIS techniques. This will become increasingly important as the amount of data grows. Moreover, because GIS is the standard, any future commercial off-the-shelf AI/ML tools for working with “dot-ology” will likely be designed for GIS.

Video-Editing Tools

Unlike FMV, which is currently run either on the Multi-INT Archive and Analysis System (MAAS)¹⁶² or Advanced Intelligence Multimedia Exploitation Suite (AIMES),¹⁶³ the WAMI architecture is newer and less-frequently used. But although the Air Force continues to invest in custom video analysis suites, a clear standard has emerged over the past few years for video editing in the commercial television and entertainment industries.¹⁶⁴ This standard is Avid,¹⁶⁵ a household name in film and television.¹⁶⁶ Avid Technology has various multimedia editing and manipulating products, and its storage system is Avid NEXIS.¹⁶⁷ Like GIS, Avid is arguably a more powerful tool than the AF DCGS really needs for current operations, and we cannot recommend that the Air Force switch to Avid at this point because we have not assessed the alternatives. We recommend that the Air Force perform the appropriate risk-benefit analysis for the use of this product. Some commercial products have been adapted well to accommodate the Air Force, such as Google Maps. More-recent experience in developing Project Maven with Google, however—in which Google canceled the contract over employee backlash—may give pause when contemplating contracting outside the more traditional defense sector.¹⁶⁸

¹⁶¹ For example, see the discussion in Environmental Systems Research Institute, “At the NGA, GIS Underpins Virtually Everything,” *ArcNews*, Spring 2017.

¹⁶² General Dynamics, “Multi-INT Analysis & Archive System (MAAS),” webpage, undated.

¹⁶³ leidos via PRNewswire, “SAIC Launches Advanced Intelligence Multimedia Exploitation Suite (AIMES),” press release, McLean, Va., November 1, 2010.

¹⁶⁴ This was not true five years ago. See Alkire et al., 2016.

¹⁶⁵ For example, every movie nominated for an Academy Award for Best Picture or Best Editing in 2018 was edited using Avid products (Igor Torgeson, “Editing Like an Oscar Winner: Why Learn Avid Media Composer?” blog post, New York Film Academy, March 5, 2018).

¹⁶⁶ Nick Messitte, “How Avid Hopes to Fix a Broken Music Industry,” *Forbes*, April 30, 2015.

¹⁶⁷ This new name removes the previous unfortunate choice from 2005, which was ISIS [the Islamic State of Iraq and Syria].

¹⁶⁸ See, for example, Daisuke Wakabayashi and Scott Shane, “Google Will Not Renew Pentagon Contract That Upset Employees,” *New York Times*, June 1, 2018.

Mission-Management Systems

AF DCGS can also make short-term improvements to the mission-management aspects of GEOINT analysis. The bent line labeled “mission metadata” on the lower left of Figure 5.1 represents a way of pulling information about the mission itself automatically from the mission management or scheduler tool so that it need not be manually entered by analysts. The “Mission Management” tool that would feed (and ultimately be fed by) GEOART today would be the UNICORN. GEOART would replace the high-altitude toolkit and FMV toolkit used today.

UNICORN is an excellent example of Air Force innovation: It was developed by airmen in response to a clear need for a tool to manage missions and disseminate GEOINT reporting.¹⁶⁹ However, as needs have arisen, UNICORN has been pressed into service to perform new functions for which it was not originally designed. Its portfolio now includes four functions: managing missions, supporting ISR assessment, facilitating reporting, and facilitating dissemination. These have been added to UNICORN largely because, as an AF DCGS–owned system, such changes could be made quickly and easily.

Through GEOART, we propose to transfer the reporting and dissemination functions to a separate AF DCGS–owned tool that could be updated separately as needed on a different development cycle to avoid overloading UNICORN and to allow for AI/ML functions to be introduced more seamlessly into the GEOINT workflow in years to come.¹⁷⁰ This should also make it easier to create a family of custom products for each CCMD or to support unclassified distribution in an HA/DR or related situation.

Taking Advantage of Future Technology to Improve GEOINT Analysis

The above short-term recommendations would help set the conditions to integrate future AI/ML into the GEOINT workflow. Generally speaking, the application of AI/ML to basic exploitation tasks would enable human effort to shift toward more-complex analysis tasks; exploiters could do more work further down the spectrum of synthesis discussed in Chapter 2. This shift would also pay dividends for future (later-phase) analysis, such as reducing the need for target systems analysts to do basic exploitation so that they can provide higher-quality products. It could also automate the population of databases, which would increase efficiency and reduce error. Table 5.2 reproduces the longer-term GEOINT recommendations from Volume 1.

¹⁶⁹ SrA Zane Wright developed the original program at the AF DCGS. See Alkire et al., 2016.

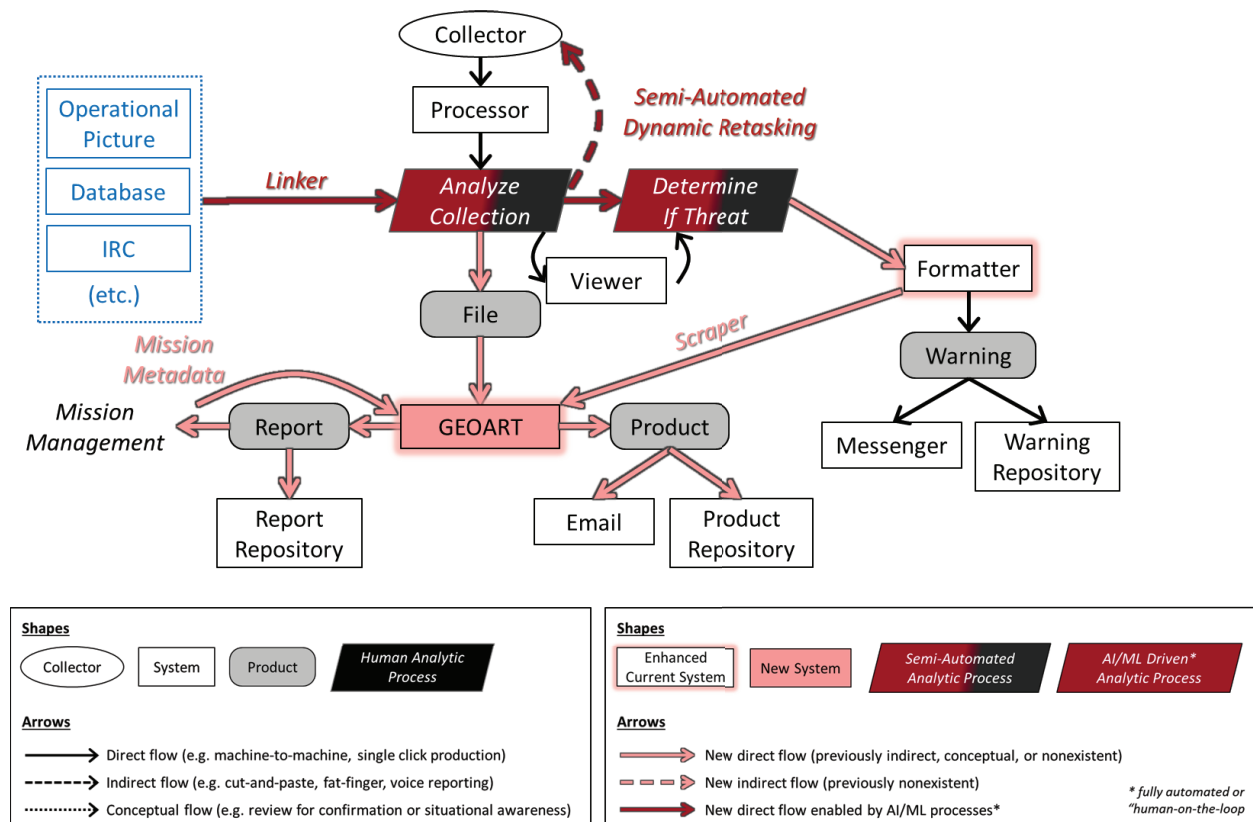
¹⁷⁰ UNICORN is in many ways the heart of AF DCGS GEOINT operations, and making upgrades while the AF DCGS is operating 24/7 is like trying to perform open-heart surgery while the patient is still awake. UNICORN has been upgraded successfully many times, but the coming AI/ML changes will pose serious design considerations. We believe that it is wiser to separate the development of these tools to ensure the AF DCGS can function more smoothly during the upgrade process.

Table 5.2. GEOINT Recommendations: Taking Advantage of Future Technology

Category	Summary of Recommendation	Analysis Roles Affected	Objectives Addressed
GEOINT	Use AI/ML to perform Phase 0 analysis and dynamically retask collectors with human-in-the-loop (or on-the-loop, depending on how the mission control element evolves).	Exploiter	Efficiency Effectiveness
	Use AI/ML to perform partial Phase 1 analysis (e.g., tag imagery, identify objects and people).	Exploiter	Efficiency Effectiveness
	Use AI/ML to flag analysts of significant changes in activity and generate threat warnings with human-on-the-loop.	Exploiter Reporter	Efficiency Effectiveness Human capital
	Use AI/ML to analyze WAMI imagery.	Exploiter	Efficiency Effectiveness
	Use AI/ML to pore through archived FMV and WAMI footage to apply partial first-phase exploitation to all and create data sets that are sharable with the IC.	Investigator	Agility
	Seek to lift the eyes-on requirement for FMV missions in which the “ISR role” indicates no risk of troops in contact (TIC) or strike decisions and when AI/ML tools can alert analysts to other events that require real-time judgments.	Exploiter	Efficiency Human capital

Figure 5.2 shows a notional data flow map indicating how future AI/ML tools would fit into the generic GEOINT workflow. Although these data flow maps must become increasingly notional as we move into the future, they provide a useful visualization of where AI/ML might be employed. In this chart, which builds on Figure 5.1, the dark red items represent the longer-term recommendations. The following sections describe these recommendations in greater detail.

Figure 5.2. Generic Data Flow Map—Potential Improvements with Future Technology



Increased Automation of Key Processes

Compared with Figure 5.1, Figure 5.2 shows a more complex process involving AI/ML and human analysts (both on-the-loop and in-the-loop). Several new items appear here. Two analytic processes, “Analyze Collection” and “Determine If Threat,” are enabled by AI/ML, and the dynamic tasking process is therefore also semiautomated. In an important change, the human-readable “Viewer” program has moved out of the main pathway and is used by the analysts only as needed. The linker has become fully automated, pulling from reference material as required. Here, we see some of the downstream dividends of integrating AI/ML tools: By populating and tagging data automatically on the back end, we can now more easily find and link to that information on the front end.

Moving the “Viewer” program reflects a larger change in the workflow. Immediately after processing, the AI/ML system would review the information. If the collection is unusable, the system would automatically request a new collection. Thus, we envision significant automation of Phase 0 tasks (although humans would need to be involved as pilots and sensor operators until AI/ML processes can be added to RPAs). The AI/ML system would then apply basic annotation and metadata, identifying as many buildings, vehicles, and other objects as it can in the image. This is Phase 1 analysis (as defined in Chapter 1). This function will not be fully automated,

however, because AI/ML is unlikely to be able to interpret human behavior as well as a human analyst can in all cases. Thus, after performing its initial analysis, the AI/ML system would flag certain collections for review by the human analyst who oversees the process and flags possible threats for immediate consideration. The human would use the “Viewer” to assess the information and may agree to issue the threat warning. For FMV collections, the analyst would also interpret human behavior and make judgments that may affect the course of the mission. For example, if the image shows two moving targets, the analyst may be able to advise the more-important one to follow. Finally, the files created by the AI/ML tool would be loaded into an enhanced GEOART that can create products or reports automatically, with the human on-the-loop to provide quality control. This would free the human analyst to review contextual material and make more-sophisticated judgment calls as needed.

Enabling AI/ML Methods and Hardware

The next set of recommendations concerns enabling capabilities. As discussed in Chapter 4, AI/ML algorithms are only as good as the data sets on which they are trained. New target classes and new environments will demand new training. Humans will therefore need to remain in-the-loop to ensure agility. The Air Force should probably lead this effort, given the uniqueness of Air Force airborne imagery in terms of viewing angle, resolution, and target set. For machine-vision applications, CNN (or progeny thereof) would be most applicable.

Given that speed of identification is not critical for still imagery—most such requirements need not be executed in near-real-time—we do not anticipate that custom hardware will be required for this particular task. For FMV, however, advanced hardware of the type discussed in Chapter 4 may be required to enable identification by AI/ML algorithms in near-real-time. Although the data sets required for Air Force applications might be unique, some of the algorithms and processes may be common to other parts of the IC. The AF DCGS should therefore look to AI/ML development efforts from the wider IC that could be leveraged for this purpose.

Finally, the same techniques used for FMV should be used to analyze WAMI imagery and pore through archival footage. The resolution differences between WAMI and FMV will require separate sets of training data, however. CNN techniques are the most likely to be used for this purpose.

Seeking to Lift the “Eyes-On” Requirement for FMV

Another potential benefit of mature AI/ML technology pertains to FMV specifically. For nearly two decades, AF DCGS analysts have labored under the requirement that at least one person in the PED crew must maintain eyes on the FMV feed at all times. Usually two or three airmen trade off this responsibility over the course of a shift. Senior Air Force leaders have expressed that this requirement adds little value. In 2011, former Air Force Vice Chief of Staff Gen James E. Cartwright called the practice: “Death TV for hours on end. It’s just a waste of

manpower.”¹⁷¹ The eyes-on requirement is not only inefficient, it also takes a significant toll on human capital. In our many interviews with airmen over the years, watching FMV feeds stands out as the most-loathed duty. Nevertheless, it is clear that airmen sometimes *should* watch certain FMV feeds in real time. There needs to be a clearer determination of when this is truly necessary.

In 2012, PAF found that asking analysts to watch more than one screen at a time would reduce their effectiveness because significant events would be more likely to be missed.¹⁷² But if the ISR mission is such that it entails no risk of TIC and no possibility of strike decision—and if AI/ML applications can alert analysts to any other events that might require real-time attention—then this concern abates. In 2015, PAF recommended that the Air Force engage with the other services to seek to relax the eyes-on requirement under such circumstances. However, at the time, there was no widely accepted method for the supported unit to designate a mission as “safe” in this regard, and AI/ML applications were not yet ready to provide backup.

A recent effort to mark appropriate missions with an ISR role has made this proposal more feasible.¹⁷³ Some of these ISR roles entail no risk of TIC and no possibility of strike decisions (such as observing patterns of life) and only infrequently require the human analyst to make real-time judgments that could affect the course of the mission (such as by redirecting the camera). When AI/ML applications reach the milestone of reliability to the point where they can detect significant events as well as analysts can, they will be ready to take over the eyes-on role for those missions. At that point, we would recommend lifting the eyes-on requirement for FMV missions in which the ISR role designated by the requestor indicates no risk of TIC and no possibility of strike decisions and when AI/ML tools can reliably flag analysts to other events that would require real-time judgment.

Considerable manpower remains tied up in FMV support. As a practical matter, the future development of the AF DCGS hinges on the ability to liberate FMV exploiters from this burden and free them up to perform tasks that make better use of human capital.

Implementing GEOINT Recommendations

As noted in Volume 1 and in Chapter 8 of this report, the *way* new tools and technologies are rolled out to users matters. For example, a recent upgrade to the GEOINT workflow—eliminating an old program that had been used to create file folder structures—met with user resistance from some senior analysts because of perceived risk. One told us that hours of work had been lost because, when the analysis program crashed, the analyst could not (or perhaps did

¹⁷¹ Ellen Nakashima and Craig Whitlock, “With Air Force’s Gorgon Drone ‘We Can See Everything,’” *Washington Post*, January 2, 2011.

¹⁷² See Menthe et al., 2012.

¹⁷³ Tingstad et al., forthcoming.

not know how to) retrieve partial products from the new hidden folder structure. We believe that better user engagement to deal with crash situations could have avoided this.

To effectively implement these GEOINT recommendations, the development team should visit (and likely revisit) multiple AF DCGS sites, including both active-duty and ANG, plus sites with special needs, such as DGS-3, to consult with the *actual analysts* who will have to use these new tools as part of their daily workflows. Reviewing training documents alone is insufficient. These documents do not explain how analysts react to all contingencies (and they could not), and much of AF DCGS crew teamwork consists of informal practices learned on the job. As discussed in Chapter 8, when the development cycle comes closer to the launch, the implementation team should identify potential opinion leaders and champions at each site to pave the way for adoption. The ultimate goal is for AI/ML technologies to help analysts do their jobs better, not to create new obstacles and irritants.

6. Rebalancing AF DCGS Competencies and Organization: Additional Detail

Chapter 2 defines the three AF DCGS competencies as (1) supporting missions, (2) supporting analysis, and (3) solving intelligence problems. AI/ML has the potential to transform each competency either directly by performing or helping to perform specific tasks associated with conducting analysis¹⁷⁴ or indirectly by freeing human analysts to focus more on solving intelligence problems and/or developing innovative ways to support and improve analysis. As the AF DCGS incorporates the AI/ML technologies recommended in Volume 1, it has an opportunity to rebalance its weight of effort among these three competencies. This chapter discusses potential future paths for the AF DCGS and their implications for organization and data flow. The discussion provides background for the “organization” recommendations summarized in Volume 1.¹⁷⁵

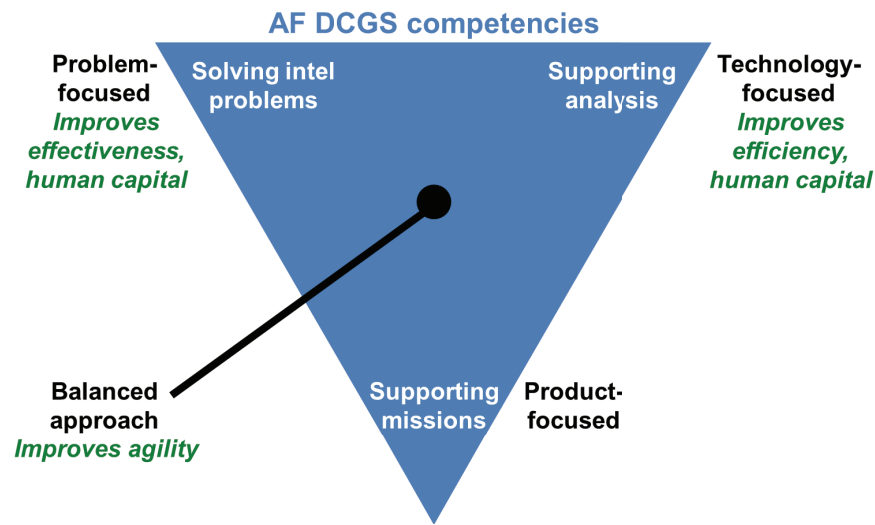
A Choice of Paths

Figure 6.1 reproduces the scheme of competencies shown in Chapter 2. The green labels indicate how shifting the weight of effort toward specific competencies would help address the challenges of efficiency, effectiveness, human capital, and agility discussed in Volume 1.

¹⁷⁴ It is important to remember that some amount of human effort will remain necessary to conduct analysis. As we discuss in Chapter 5, even after leveraging AI/ML methods for key tasks, there will be many tasks where data flow not only benefits from but requires human intervention.

¹⁷⁵ Menthe et al., 2021.

Figure 6.1. Rebalancing Weight of Effort Across AF DCGS Competencies



Solving Intelligence Problems

One possible direction would be to shift the weight of human effort toward solving intelligence problems. This implies a *problem-focused* organization, where the aim is to push human effort down the spectrum of synthesis (discussed in Chapter 2) toward all-source intelligence and multi-INT fusion. In this construct, analysts would no longer need to be organized in *crews*—a structure that derives from the idea that each intelligence-collection platform needs a dedicated group of people to exploit information during or after every mission. The problems that each team focuses on would instead depend on joint, IC, service, and air component priorities. A team could also simply be assigned to support a given ground unit or tactical area of operations.¹⁷⁶

With analytic work separated from specific platforms via automation, there would have to be a separate mechanism (in addition to changing the Global Force Management process’s implicit assumption that AF DCGS operations be organized by platform) for apportioning analyst time to specific problems. It will also be important to consider the unique value that AF DCGS analysts offer to different needs, such as threat warning for aircraft and tactical support to ground, amphibious, and maritime forces. Some of the more unusual items might need separate “problem” teams. Shifting AF DCGS weight of effort toward this competency would help improve overall effectiveness at supporting warfighters and make better use of the skills that human (versus artificial) intelligence brings to the table.

¹⁷⁶ PAF has suggested an area-centric scheme for conducting PED for some time. See Menthe et al., 2012.

Supporting Analysis

Another option would be to shift toward supporting analysis, an approach that would become largely *technology focused* as AI/ML tools are introduced. As we discuss in Chapter 8, developing, onboarding, maintaining, and evolving new technologies for the AF DCGS requires an immense amount of vigilance, analysis, and expertise—not just in programming or the *how* that supports innovation but also in the *what* and the *why* of the missions the technology is helping to achieve. Technology needs humans to work. For at least the next several years or decades, humans must define and evolve requirements, test the fidelity and usefulness of tools, catch mistakes, conceive of new applications for tools, provide judgment and intuition, and fulfill reasonable legal requirements (e.g., laws prohibiting machines from standing alone in the kill chain).

Analysts within the AF DCGS (and elsewhere) are already doing parts of this work, typically on an informal or ad hoc basis. For example, analysts might decide to fix a technological glitch on their own time. They might participate in crowd-sourcing activities to support a ML project. One way to formalize this effort would be to align sites or teams using the data format (e.g., still imagery, text, voice, signal) that automated tools are attempting to exploit or analyze. There might be additional groups that are dedicated to integrating multiple intelligence sources and ensuring the flow of information among databases and between databases and users. The objective would be to improve the efficiency of AF DCGS analysis, whether human or machine.¹⁷⁷

Balancing for Agility

The best way forward would be for the AF DCGS to balance the approaches just mentioned. Including all three competencies in future AF DCGS plans will keep the organization agile and responsive to evolving national security needs. Of course, retaining the three competencies in one organization will require the identification and delineation of teams across the enterprise to work on each. Individual sites need not include all three but could do so. For example, some locations might need to provide greater support to a theater's problem-centric needs, whereas others might place greater emphasis on being a testbed for technological innovation.

The ratio of emphasis among the competencies can—and should—vary over time. In the coming years, as technology comes on board, DGS locations will likely still maintain large numbers of teams on the ops floor that are focused on supporting missions, with smaller teams

¹⁷⁷ We note that increasing technological specialization may also drive toward more-complex skill sets. As a practical matter, this may lead to further subdividing of skills as the INTs become increasingly specialized, much as technical electronic intelligence has become a subset of SIGINT. For example: “As next-generation FMV and still imagery sensor technology are fielded, the two jobs will become more distinct and require dedicated skill sets” (Jennifer A. Hollock, *Flexibility Versus Expertise: A Closer Look at the Employment of United States Air Force Imagery Analysts*, master's thesis, Maxwell Air Force Base, Ala.: Air Command And Staff College, Air University, October 2017, p. 2).

focused on the other two competencies. Over time, as more steps in the process become automated, the number of teams focused on supporting missions might decrease in favor of supporting analysis improvements and solving intelligence problems. The role of supporting analysis will shift toward the ops floor as analysts are freed from the burden of performing more-basic tasks and can consider how to improve those processes. The DART will increasingly focus on intelligence questions and will likely grow.

It is important that the AF DCGS strive to balance all three competencies, even when there is a temporary shortage of demand for one or more of them. Otherwise, the AF DCGS could devolve into supporting only one of the competencies, even though all are needed in the long term. There are, unfortunately, plenty of examples throughout history where analytic capabilities were suspended, only to become important some years later, when personnel with relevant skills were no longer available.¹⁷⁸ Thus, the AF DCGS would do well to cover its bases and ensure that some level of analytic capability exists within all three competencies, ready to surge if and when needed.

Organizing Around Phases for Scalability

Another way to think about the rebalancing of AF DCGS competencies is to envision different parts of the AF DCGS as organized around different phases of analysis. In time, AI/ML tools will become increasingly adept at performing Phase 0 analysis (determining usefulness of collections, dynamically retasking to retake a failed collection, issuing threat warnings) and parts of Phase 1 analysis (answering basic “who,” “what,” “when,” and “where” questions) for intelligence as it comes in. When they can do so reliably—achieving success and false alarm rates at least comparable to human analysts—then analysts will be able to focus almost exclusively on the other parts of Phase 1 analysis (answering more-advanced “why” and “how” questions) and on conducting Phase 1.5 analysis and beyond (correlating multisource, multi-INT information to solve intelligence problems as needed).

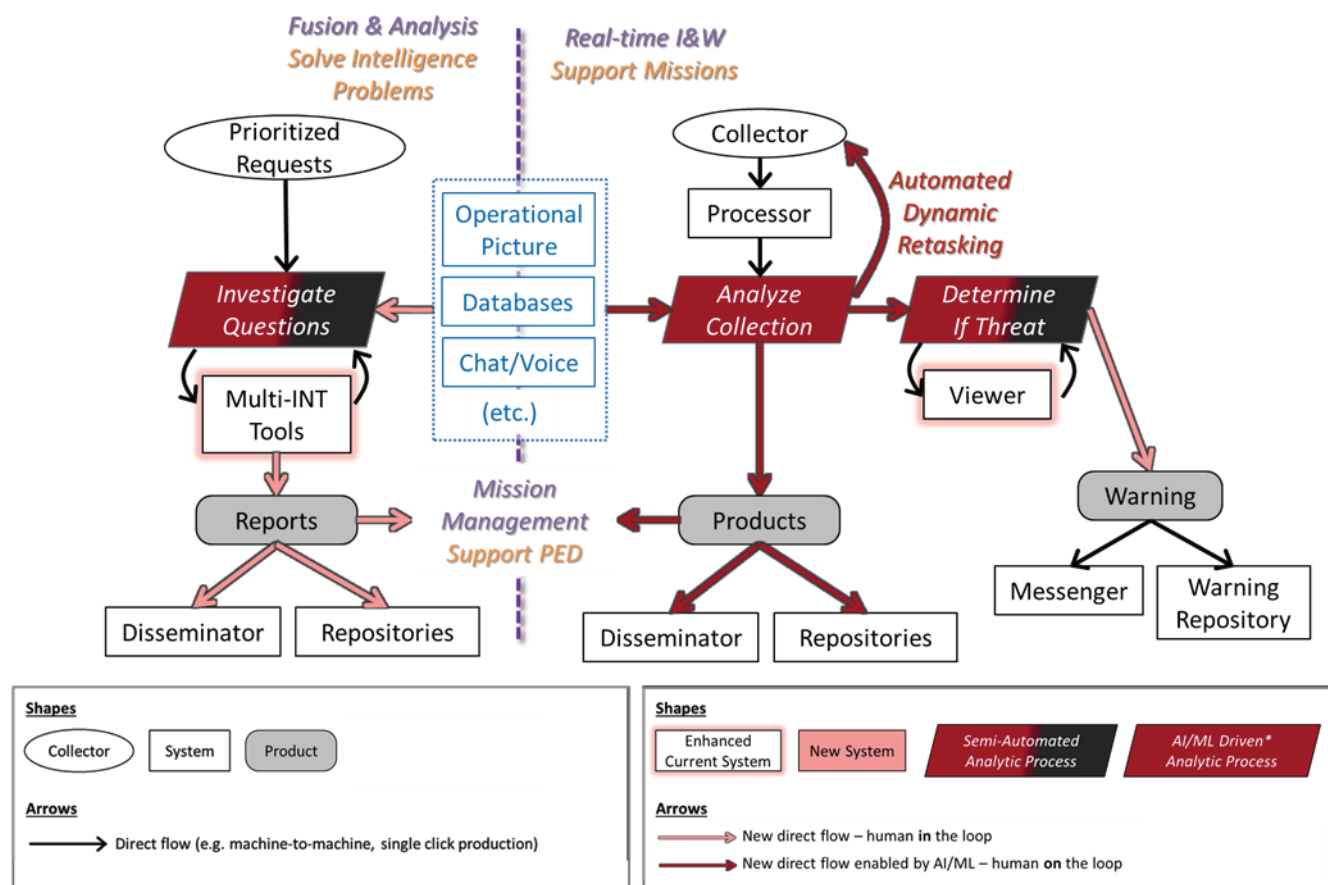
By organizing machine and human resources around different phases, the AF DCGS can make better use of human capital, improve efficiency and effectiveness for both human and machine tasks, and allow more-agile scaling of resources as needs evolve. AF DCGS can standardize the use of AI/ML for Phase 0 and parts of Phase 1 analysis for all collections as soon as they are ingested to create a tagged, searchable database. Human analysts can thus be freed from specific input sources (performing more-advanced parts of Phase 1 analysis only when needed to answer specific questions) and can focus instead on the desired output—the

¹⁷⁸ Some argue, for example, that the Air Force’s expertise in conducting irregular warfare during the Vietnam era was largely lost and had to be reconstituted in the 21st century (see Will Sellber, “The Other Side of the COIN,” *Air and Space Power Journal*, Vol. 32, No. 3, Fall 2018).

customer's needs. We may think of this as a shift in human effort from input-centric, “as-collected” analysis to output-centric, “on-demand” analysis.¹⁷⁹

This shift suggests a corresponding change in AF DCGS organization. As illustrated in the generic data flow map in Figure 6.2, we may envision a future state in which each of the competencies just discussed is addressed by a different team. This is designed to grow organically, we hope, out of the current ops-floor/DART split.

Figure 6.2. Generic Data Flow Map—Potential Organization to Balance Competencies



NOTE: I&W = indication and warning. * = fully automatic or human-on-the-loop.

In this schematic, the basic questions (who, what, when, and where) would be answered by AI/ML for all collections as they arrive.¹⁸⁰ The advanced questions (how and why) would be

¹⁷⁹ Although PAF has discussed this shift in earlier reports (e.g., Alkire et al., 2016), we are not alone. As a previous 480th ISRW commander observed: “Ideally, analytical elements such as AF DCGS should not ‘chase’ airborne ISR collection but instead should analyze and exploit any and all sources available that will successfully answer the questions posed by the supported commander” (Haugh and Leonard, 2017, p. 11).

¹⁸⁰ As noted in Chapter 5, there may be something of a “PED effect” like the “AI effect,” where, once computers are able to perform these tasks, they are no longer thought of as PED.

answered by humans, but only as needed. Human analysts would answer intelligence questions, fulfill requests for information, support a ground unit, or warn of threats in a human-on-the-loop fashion. When they need to perform additional analysis, human analysts would sift through AI/ML-tagged data in the cloud rather than manipulate data as they come in from the sensors. Mission management would oversee both halves of the construct, tracking progress and developing tools and training as appropriate to each.

Instead of scaling the number of crews with the number of collectors, this construct would scale human effort with the number of supported units or requestors of information. This is a change from an input-scaled to an output-scaled structure. It is also a shift of human effort from analyzing everything “as collected” to providing additional analysis “on-demand.” This is a more sustainable and scalable approach than attempting to conduct full analysis for all collected data, and it also makes better use of human capital by moving human effort down the spectrum of synthesis toward multi-INT and all-source analysis, as the new SIAS paradigm anticipates.

7. Building the Right Skills: Additional Detail

Although much attention is devoted to how technological innovation can improve AF DCGS functions, teams of well-trained human analysts remain the true heart of the enterprise. As Lt Col Jason M. Brown wrote in 2009: “the quality of the DCGS is defined less by machines and more by the complex and largely intangible web of human behaviors and abilities—the *human factor* within the system.”¹⁸¹ This has been echoed in PAF’s previous work in this area: “When it comes to PED, the Air Force’s most valuable asset is, and will remain, its force of trained human analysts”¹⁸² and “[w]hen it comes to tools and technologies, empowering [the Air Force’s] analysts should be the heart of USAF’s [the U.S. Air Force’s] analytic strategy.”¹⁸³ As the AF DCGS moves to introduce greater automation into its processes, two specific human contributions stand out as important: to do what automation cannot yet do and to compensate for vulnerabilities created or exacerbated by automation.

Project Maven, the current DoD effort to train AI/ML to classify objects in FMV, illustrates both the benefits and deficiencies of automating individual steps in the analysis process. On the one hand, the capability to automatically and accurately classify an object as a truck can save analysts time and help with cataloging archived imagery. On the other hand, object classification is only one part of a longer analytic process. Many subsequent steps are needed to provide key context, such as identifying *which* truck is in the image, *who* it belongs to, and *where* it has been observed before. As with all processes, automating one step but not others may simply shift the bottleneck elsewhere.

Moreover, exploitation is not as basic as it seems. Analysts are increasingly asked to interpret the behavior of adversaries and civilian populations. Distinguishing between friend and foe and understanding social relationships can be complex determinations requiring human judgment, cultural knowledge, and other skills. These analytic processes are, of course, more varied and difficult to automate.

There is also increasing recognition that humans may need to be trained so that they can step in if automation fails. The U.S. Navy provides one case in point. The U.S. Naval Academy phased out its celestial navigation classes in 2006, but it brought back the venerable sextant a decade later because of concerns that absolute reliance on the Global Positioning System (GPS) created a vulnerability in the event of cyberattacks, GPS jamming, or orbital ablation of

¹⁸¹ Jason M. Brown, “Operating the Distributed Common Ground System: A Look at the Human Factor in Net-Centric Operations,” *Air and Space Power Journal*, Winter 2009, p. 52.

¹⁸² Menthe, Cordova, et al., 2015, p. 57.

¹⁸³ Alkire et al., 2016.

satellites.¹⁸⁴ As we discussed at the end of Chapter 4, the use of AI/ML algorithms for intelligence analysis also creates new vulnerabilities that might require human mitigation.

A more technologically enabled AF DCGS, such as what is envisioned in this research, would also require analysts to have a level of technological proficiency that is appropriate to their missions and consistent with the tools they must use or supervise. In some cases, analysts may already have a basic or even advanced proficiency because of their use of technology in their personal lives or prior job or educational experience—and we have seen examples of this.¹⁸⁵ In other cases, skills may need to be taught.

Volume 1 outlines steps that the AF DCGS should take to build the human skills needed to ensure proper functioning of AI/ML and to take advantage of the opportunities of AI/ML.¹⁸⁶ This chapter provides additional context for and discussion of those recommendations. Many of the ideas to expand existing or trial AF DCGS programs arose during conversations with the analysts implementing them.

Building Programming and Data Science Skills

Any infusion of tools that address aspects of basic exploitation and reporting and make data more accessible for analysts will decrease the demand for skills devoted to single-INT exploitation. However, these proficiencies must be maintained on some level so that tools can continue to be monitored and improved on by humans. We therefore recommend that the AF DCGS retain basic INT skills for some airmen, even where they may appear obsolete because of AI/ML.

But these tools will also present an opportunity for AF DCGS airmen to regularly engage in more-complex analysis, which could include multiple intelligence sources, developing deep understanding of a target set over time, and/or working on the cutting edge of tool or data science technique development. This will necessitate developing or bringing on airmen with somewhat more analytically focused skills than what we see today.

The rapidly changing technological environment points to the need to train young enlisted analysts to thrive in a changing work environment and the challenge to retaining them. The junior enlisted airmen in the Air Force ISR community represent a treasure trove of energy that has yet to be fully tapped. During our research, we encountered many who wanted to learn and make a difference. Air Force ISR leaders should consider exposing junior analysts to more-advanced techniques soon after they attain initial competence in their individual intelligence discipline.

¹⁸⁴ David Dickinson, “Navy Resumes Celestial Navigation Course,” *Sky and Telescope*, April 5, 2016.

¹⁸⁵ Unfortunately, this also means that analysts may be increasingly skeptical or dismissive of tools that do not work as well as those to which they have become accustomed in their civilian lives. Including analysts early in the innovation process is crucial, as discussed in Chapter 8.

¹⁸⁶ Menthe et al., 2021.

All analysts will need consistent proficiency in technologies to allow them to team up with machines as appropriate. Regardless of how the AF DCGS evolves, it will be necessary for all analysts to be comfortable using tools and have some analysts with basic knowledge (e.g., two or three college courses) in programming and/or data science to help monitor how well their tools are working, identify opportunities for growth, and communicate problems and new requirements.

Providing analysts with opportunities to better understand the promise of all-source analysis (including which databases support it) will be important as they form problem-centric teams, regardless of the AF DCGS competencies in which they primarily work. Whether these airmen are supporting missions or attempting to solve more-complex intelligence problems, a basic understanding of data-science techniques will be increasingly useful.

One specific capability that will be needed is the ability to work with a GIS, such as ArcGIS, as discussed in Chapter 5. For analysts working to support analysis, a basic understanding of programming will also be needed to appreciate the opportunities and limits of algorithms and to communicate these needs to professional programmers.¹⁸⁷ Some of this training in data-science techniques (such as GIS basics) should probably be delivered at Goodfellow AFB along with basic training. Additional emphasis on embedded or on-the-job training is also worth considering. Many analyst airmen reported to us that, beyond the most basic training required to log on to their computer systems, there is simply no substitute for on-the-job experience.

If automation could smooth the workflow and handle repetitive, time-consuming tasks, then senior analyst airmen would be able to take advantage of training opportunities more consistently. It may also be valuable to take advantage of the copious university and college certificate programs in AI, programming, data science, GIS, and other such techniques. Analysts often must wait months to start “real” mission work after basic training, usually because of the glacial pace of adjudicating applications for certain security clearances. This could be the perfect time to take a certificate course. Another option might be to use the opportunity between assignments as educational opportunities.

Certificates provide background and basic competency without overemphasis. We do not foresee that the Air Force needs a large cadre of programmers and data scientists, at least not one that is likely to grow substantially beyond what is available now. This is because the Air Force will likely never be able to peel away the resources to support the kind of cadre that commercial industries can. Thus, it makes sense, in many cases, for Air Force analysts to be knowledgeable enough to converse, but not directly compete, with civilian programming experts.

One training option, in addition to changes in basic training and embedded training, could be to extend the current Combat Readiness Sustainment Program that was introduced in the latter half of 2017 for SIGINT crews and is now being used more widely to provide opportunities for

¹⁸⁷ Technical literacy is a lesser requirement than technical mastery but still requires some basic courses and hands-on training.

mid-career add-on training.¹⁸⁸ This program essentially allows analysts to take “sabbaticals” of different lengths (depending on their tenure within the AF DCGS) to pursue different skills. This allows some training to be tailored to the needs of different analysts and sites.

Sharing and Capturing Tactics, Techniques, and Procedures

As the Air Force ISR community continues to develop and deploy new sensors, if it is to find new ways to use existing sensors and apply improved analytic tools to the data it collects, two TTP-related issues often arise: sharing and capturing.

The first issue is the challenge of *sharing* a new method or best practice that has been developed to leverage a new capability. One analyst or one site may have invented a new method or adopted a great tool, but complexity, lack of communication, and even lack of trust can hinder wider implementation across the AF DCGS. ISR WEPTAC conferences can be an effective means of sharing new TTPs if the conferences are sufficiently frequent and geared to the tactical level, which is not always the case.

The second issue is the problem of *capturing* the necessary analytic TTP associated with new capabilities for the long term. Intelligence analysis is not a simple process, and, at the AF DCGS—as in any technical organization—many of the little details or “tricks” needed to work with certain systems or manipulate complex data formats are not documented but are learned by experience or passed down by word of mouth. When the mission is routine, this works well, but when the mission involves employing a niche capability, having knowledge of how to analyze data to support that mission could be fleeting. When the specific situation that gave rise to the need for a novel method ends and that method ceases to be employed on a regular basis, such knowledge can atrophy or be lost altogether. When similar needs arise in the future, analysts experience *déjà vu* as they attempt to reinvent the wheel.

Reaching the full potential of a new method, system, or technology requires standardizing TTPs for employment across the ISR enterprise and the operational force. Furthermore, fixing issues requires a concerted and holistic approach with tracking and follow-up to ensure that initiatives are resourced to completion. New methods need to be codified in such a way that if they are not used because of changing requirements, they can be quickly brought back off the shelf if needed again later.

Rehearsal-of-Concept Drills

Another promising vehicle for building airmen’s skills as technologies evolve is the rehearsal-of-concept (RoC) drills that began in the AF DCGS in late 2017. RoC drills are DGS and DMS training events that are conducted every two or three months to train airmen in their

¹⁸⁸ 480th Intelligence, Surveillance and Reconnaissance Wing, “480th ISRW Institutes Combat Readiness Sustainment Program,” Air Combat Command webpage, October 30, 2017.

roles, tools, processes, and team collaboration for a given vignette. They are also used as part of the Combat Readiness Sustainment Program. RoC drills help airmen maintain proficiency across current mission sets and major operations scenarios that are not practiced in day-to-day operations. They usually last for a day or half a day.

The drills occur as a tabletop exercise rather than the type of operational exercise that is required during a standard certification process. They provide an environment for all levels (E-1 to E-7) to learn. They are meant to develop critical thinking and create an understanding across the ops floor of what different analysts are doing and how they contribute to the big picture. They also involve regular “interruptions” or periodic evaluation sessions in which the participants can reflect on their progress so far in the drill. RoC events to date have increased awareness between the GEOINT and SIGINT staffs of one another’s capabilities and jobs. Because they require minimal use of computers and other aides, the only real cost is scenario development and execution time.

One thing that makes RoCs unique is that they are designed as *puzzle-solving* exercises. A situation is presented that no single analyst or PED crew can solve on their own. Success requires discussions with other PED crews, learning their capabilities, and putting together a collaborative plan of action for coordinating activity across the ops floor. As might be imagined, designing an exercise that works in this manner is difficult, and not every instance will succeed. Moderating such an exercise also takes skill. We believe that enlisted airmen have these skills.¹⁸⁹

We hypothesize that the RoC events can help the AF DCGS examine new tools, skill sets, and processes that can help address existing and emerging analytic challenges. A tailored RoC event could examine a particular challenge in a scenario context to see what analytic roles and capabilities are needed, can be improved on, or may emerge. The goal would be to structure an RoC drill around such challenges as

- flooded data environment
- sparse data environment
- rapid increase in operational tempo
- need to rapidly share information with new partners
- emergence of a new focus area or information source requiring updates to TTP
- integration of a new intelligence partner whose TTP and products need to be understood.

The event would leverage the expertise of the AF DCGS analysts and provide another forum for exploring innovative analysis. For example, imagine a drill involving a time-sensitive target mission, where the critical metric is speed of engagement. The team would be briefed to the mission and would take their seats around the table. The scenario would “run” for a period of

¹⁸⁹ As one airman explained, creating and running a RoC drill draws on skills that are similar to being a Game Master in a Dungeons and Dragons role-playing campaign. It is worth noting that, in an unrelated context, one DGS commander suggested that the Air Force recruit analysts for the AF DCGS at tabletop game conventions rather than “just monster truck rallies” because they needed puzzle-solvers. As the weight of human effort in the AF DCGS shifts more toward the investigator role, these insights may be of interest to the Air Force.

time, then freeze. During the “run,” players would perform their jobs (tactical level). They would respond to events on the ground, analyze ISR data, and perhaps make decisions about using new tools, information, or partners available to them. For example, participants might decide to select among different “new enabler” choices (e.g., new collection capabilities or analytic methods) and would then need to demonstrate an ability to integrate the choice into their typical workflow. At the pause, the team would evaluate the process that occurred so far. They may answer such questions as:

- What was your mission objective in terms of timeliness?
- Were you able to respond as quickly as you needed to?
- What were the roadblocks in the process, if any?
- What additional skill set do you think is needed that you might not currently have?
- Are there improvements to the process that would help meet the timeliness objective (improvements may cover analysis tools, databases, TTP, etc.)?
- What enabled you to achieve your objective?

Although RoC drills currently address specific training objectives and are used to foster team-building, there is an opportunity to build on this mechanism to provide input for future needs or approaches to doing business. They could also serve as a means of helping inform innovation for the entire enterprise. This endeavor would have the added benefit of engaging users across the enterprise in an interactive way by including them in the concept-formulation phase of innovation. As we see in the Chapter 8, analyst engagement is critical to the success of new technologies within the AF DCGS.

Using Mission Type Orders and Focused Collections as Testbeds

As information sources and their associated databases proliferate, analysts must continue to keep up with their knowledge of what is available. They cannot use what they do not know exists. They must also learn to use fusion tools for combining this information to tease greater information out of what has been collected, especially in sensitive areas. Several such multi-INT tools are available today. Moreover, GIS tools will be increasingly important in the future. Finally, one of the most important contributions that human analysts can make is to go beyond monitoring to discovery, addressing the most bedeviling of intelligence problems: finding the “unknown unknowns.”

There is no silver bullet we can recommend here, but mission type orders (MTOs) and focused collection operations can provide a useful testbed for seeking new ways to approach difficult ISR problems. We urge the Air Force to encourage CCMDs to allow the AF DCGS to provide support via MTOs and focused collection operations. These allow analysts to become more engaged in ISR operations and to help understand the “why” of their work. Many analysts and leaders at various AF DCGS sites emphasized the value of MTOs, not just for improving collection but also for nurturing human capital.

8. Fostering Innovation and Successful Implementation: Additional Detail

Most Air Force analysts who served during the height of the Cold War grew up in an era when color televisions were new. These analysts might have pored over images developed from a U-2 wet-film optical bar camera, searching for any detail that might signify a change in the strategic winds. By the turn of the 21st century, an analyst airman might have owned a Nokia bar phone and would have entered service hearing stories about the Gulf War. These analysts would likely have focused on deliberate targeting and routine collections on fixed installations to maintain order-of-battle information on potential adversaries—and these were the airmen called on to be the first FMV analysts on the line for the long ISR surge.

Now, a new generation of analyst airmen has grown up with smartphones and wireless internet: In their everyday lives, they carry in their pockets the means to access information almost anywhere about almost anything. Many expect to see drones delivering their household supplies and self-driving cars well within their lifetimes.¹⁹⁰ This new generation of analyst airmen has learned to track fleeting targets, support TIC at a moment's notice, analyze data from advanced sensors, and work in virtual teams whose members are physically distributed around the globe. But although they regularly integrate new technologies into their personal lives today, doing so in the AF DCGS environment is not so simple.

Our observations show three reasons for this difference. First, the purpose of the work is clearly different. Ordering an item over the web or posting on social media is different from chasing a terrorist cell or issuing threat warnings for air defense because lives may depend on the outcome. Second, AF DCGS work is more structured (e.g., in its use of protocols) and requires coordination with more people than tasks in everyday life. Personal uses tend to be more self-directed and have larger room for error. Third, commercial technology has been specifically designed for easy adoption, while AF DCGS technology has not. Nevertheless, the AF DCGS can learn from best practices in other sectors to create an environment in which new tools are welcomed and have a greater chance of being adopted.

Volume 1 proposes steps that the AF DCGS can take to ensure smoother onboarding of new technologies and foster a culture of innovation.¹⁹¹ These recommendations are based on our analyses of case studies and literature on best practices for developing and introducing new technology. This chapter presents that analysis in greater detail.

¹⁹⁰ Already, millions of drones are in use. In summer 2016, nearly 75 percent of consumers surveyed expected drone-enabled deliveries to begin within the next five years (see Tom Standage, “Taking Flight,” *The Economist: Technology Quarterly*, June 8, 2017).

¹⁹¹ Menthe et al., 2021.

Confronting the Innovation Adoption Problem

An important finding from early interviews for our research is that the way in which new tools are developed and brought into the AF DCGS ecosystem can greatly influence how well they are integrated and whether they ultimately improve performance. Well-meaning developers have sometimes designed prototype tools that look great on paper but do not meet analysts' needs.¹⁹² Tools that appear on an analyst's desktop without much explanation or training are likely to be left to languish in the corner unless a patient risk-taker can identify a good use and evangelize the concept to others. These issues are hardly unique to the AF DCGS. As one expert put it: "DoD does not have an innovation problem; it has an innovation *adoption* problem."¹⁹³

To be successful, an onboarding process for new tools for the AF DCGS must be holistic: It must consider the end-to-end process of designing, introducing, training for, and sustaining those tools. It must also take into account that "newer" is not necessarily better, and that even useful tools need to be replaced when they have become stale. In this chapter, we argue—as others have—that the Air Force must embrace a different mentality for onboarding software than has historically been used to acquire hardware. Applications, digital interfaces, algorithms, and other "soft" system elements can and should be built "as the aircraft is flying." Given the need to test, fix, and retest many times, this requires a flexible environment in which to receive and act on rapid "bottom-up" feedback from discerning users.

In this chapter, however, we do not address certain development cycle issues, such as release schedules, software "beta" processes, and cost-estimation techniques. The stages of onboarding tools described here are aimed primarily at guiding smaller development efforts that can be pursued within the existing DoD acquisition framework, not larger development efforts that may benefit from reforms to the formal DoD acquisition system itself. Planning for advanced AI/ML development projects will likely entail additional consideration beyond the scope of this project. The more general problem of applying DoD processes to information technology acquisition—particularly to the rapid acquisition of technologies that may entail risk—remains an important topic of continuing research.¹⁹⁴

Next, we present one framework for understanding the process and institutional needs for effectively onboarding tools for the AF DCGS. We include case studies of innovation in the commercial sector and some examples of successful innovation in the IC.

¹⁹² As one analyst we spoke to put it plaintively, "They give us so many tools, but no one can help us."

¹⁹³ Eric Schmidt, "Statement of Dr. Eric Schmidt, House Armed Services Committee," Washington, D.C.: U.S. House of Representatives, April 17, 2018 (emphasis in original).

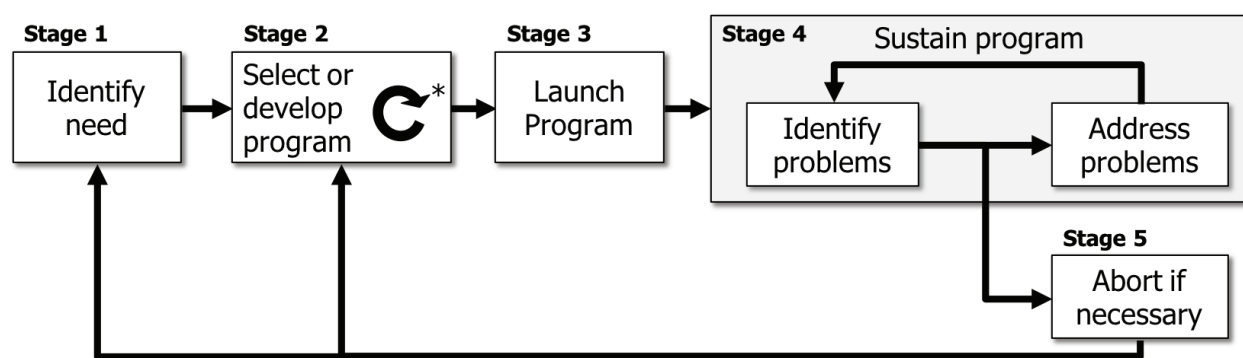
¹⁹⁴ See, for example, Jeffrey A. Drezner and Michael Simpson, *Exploring Parallel Development in the Context of Agile Acquisition: Analytical Support to the Air Superiority 2030 Enterprise Capability Collaboration Team*, Santa Monica, Calif.: RAND Corporation, RR-1808-AF, 2017; and John Birkler, *Untying Gulliver: Taking Risks to Acquire Novel Weapon Systems*, Santa Monica, Calif.: RAND Corporation, OP-268-OSD, 2009.

Stages of Technology Implementation

The literature surrounding program implementation strategies—particularly for technology—spans many fields, including behavioral science, organizational psychology, change management, and implementation science. Case studies provide lessons learned from distinct programs and approaches that can be applied to Air Force implementations. Based on our review of the literature, conversations with SMEs, and case studies both within and beyond the defense sector, we have summarized a preliminary framework of the technology implementation cycle to assist Air Force leadership in fostering innovation and onboarding new tools.

Figure 8.1 presents an overview of the five stages of the technology implementation life cycle. Those familiar with DoD acquisition may wonder why this framework looks more linear than cyclical. The simple answer is that it is more cyclical than it appears at first glance and contains several embedded development loops—this is indeed primarily an agile process. But the deeper answer is that, as a formal weapons system that thousands of airmen must be trained to use, the AF DCGS requires some measure of standardization. Although regular, even continual improvement of processes is desirable, the AF DCGS would be ill-advised to stay in a perpetual “beta” development state. As one commander, frustrated by updates and rollbacks, explained when he was asked what new tool or technology he most desired: “I just want what I have to actually work.”¹⁹⁵ It is also important to recognize that there *are* distinct steps that must be taken to shepherd an idea from conception to implementation, even though, for any given program, the transitions between these stages are not always sharp and do not follow the same prescribed timeline. This can be helpful in ensuring that people with the right skills are able to implement or champion appropriate steps.

Figure 8.1. Stages of Technology Implementation



NOTE: *The development cycle involves rapid iteration through multiple builds.

¹⁹⁵ This interview was part of a previous PAF project on PED.

Beginning with the initial development cycle, the need—the problem that the program is trying to solve—must first be identified (Stage 1). The program is then developed, either in-house or with the assistance of a contractor, or selected and procured from a vendor (Stage 2). As indicated in the diagram, there is considerable iteration and user engagement in this process. The launch of the program (in full or in pilot) across the enterprise is where widespread implementation truly begins (Stage 3). However, as the initial excitement fades,¹⁹⁶ significant effort and resources must be dedicated to sustaining the program by circumnavigating any challenges that appear in the months or years that follow (Stage 4). When these challenges become severe, leadership must decide whether to stay the course (with adjustments to correct the identified issues) or abort the program (Stage 5) and return to the drawing board. We note that all software is ultimately transient, and all tools will need to be offboarded eventually. Doing so efficiently, even ruthlessly, can be important to minimizing unnecessary training on obsolete systems.

This overview is designed to be broadly applicable but will need further refining in order to be applied to individual Air Force programs. Future studies could allow the validation of the framework and strengthen its applicability to specific areas. We now walk through the steps in more detail.

Stage 1: Identify the Need

Identifying a need is often a major challenge. Previous PAF research on innovation within the Air Force concluded, “While many believe technological change is the root cause of military innovation, our research indicates that major Air Force innovations usually start with the identification and framing of a strategically important operational problem.”¹⁹⁷ In identifying a need, it is imperative to ensure that the need is tied to the people using a software program (including the environment in which they plan to use it) and that the need is clear to the users.

Past PAF research on force modernization has emphasized that intended users or those close to the users play a crucial role in defining key operational goals—including identifying needs.¹⁹⁸ In this stage, airmen should play the lead role in determining their own needs, although contractors may serve a valuable role in facilitating these conversations.

In our interviews and observations at many DGS sites, we have encountered remarkable creativity and innovation on the part of airmen.¹⁹⁹ We have met some who voluntarily flowchart

¹⁹⁶ If there is no initial excitement, that is already a warning sign.

¹⁹⁷ Adam R. Grissom, Caitlin Lee, and Karl P. Mueller, *Innovation in the United States Air Force: Evidence from Six Cases*, Santa Monica, Calif.: RAND Corporation, RR-1207-AF, 2016, p. vii.

¹⁹⁸ David Ochmanek, “Promoting Innovation and Modernization Within the Air Force,” Santa Monica, Calif.: RAND Corporation, RB-99-AF, 2003.

¹⁹⁹ There is some selection bias here, of course, because the airmen most willing to speak to the PAF team are also those most likely to be self-motivated. But we have also spent dozens of hours walking freely about the ops floor,

and redesign their own analytic processes. We have met several who have taken the initiative to learn Python coding on their own—some by watching YouTube videos or using other internet educational resources—to write process data for web interfaces or script ArcGIS processes. We have even met AF DCGS airmen working to improve collection capabilities. The DART and the AF DCGS ops floor are fertile ground for sowing new tools, technologies, and processes.

Stage 2: Select or Develop the Program

The design and functionality of the program itself, regardless of how well its implementation is carried out, has a strong effect on how well the technology is ultimately accepted by the users and the success of the implementation as a whole. The literature has identified several (often similar) constructs that are effective in predicting technology acceptance.²⁰⁰ Two major factors are *perceived ease of use* and *perceived usefulness*.²⁰¹ Perceived ease of use involves how simple a program is to understand and use. Prioritizing the development of program characteristics, such as user interfaces, can increase the likelihood of user adoption in later stages. System characteristics, in particular, can be designed with users in mind during the development stage. This applies generally in the sense that programs should be intuitive to use and, when they are not, they can engender frustration and user resistance (discussed in more detail later). Perceived usefulness is how much an employee believes that the new program will help them improve their performance. When selecting or developing a program, perceived usefulness can be enhanced by ensuring that the program aligns to demonstrated needs and operational realities of the users and/or is an improvement over an existing program.

In addition to the obvious benefits derived from aligning user needs and product capability, integrating users into the development or selection of a program ultimately increases user satisfaction and buy-in.²⁰² Other factors, such as social influence, user training, and the implementation process, can also influence user acceptance.²⁰³ For programs with substantial AI components, it is helpful to incorporate an explainable model, paired with an explainable interface that helps the user understand why the algorithm makes the choices it does, when it succeeds or fails, and when it can be trusted.²⁰⁴

interacting with various PED crews and DART personnel, and, in our experience, airmen who wish only to complete their checklists and go home are more the exception than the rule.

²⁰⁰ Nikola Marangunić and Andrina Granić, “Technology Acceptance Model: A Literature Review from 1986 to 2013,” *Universal Access in the Information Society*, Vol. 14, No. 1, February 16, 2014.

²⁰¹ Viswanath Venkatesh, Michael G. Morris, Gordon B. Davis, and Fred D. Davis, “User Acceptance of Information Technology: Toward a Unified View,” *MIS Quarterly*, Vol. 27, No. 3, September 2003.

²⁰² Dorothy Leonard-Barton and William A. Kraus, “Implementing New Technology,” *Harvard Business Review*, November 1985.

²⁰³ Marangunić and Granić, 2014.

²⁰⁴ Defense Advanced Research Projects Agency, *Broad Agency Announcement: Explainable Artificial Intelligence (XAI)*, Arlington, Va., DARPA-BAA-16-53, August 10, 2016.

It may be beneficial to outsource development to contractors when there is insufficient manpower and/or technical expertise in-house. For the Air Force, this includes situations in which uniformed personnel have competing demands on their time that do not allow for the focused effort needed to develop the program internally. In those cases, partnering with a contractor who can provide that dedicated attention over time should be considered. However, outsourcing is not without risk, as it might result in the atrophy of in-house expertise, could make oversight of the contracted project more challenging, and could negatively affect airmen's development of technical skills in the long term.²⁰⁵ Outsourcing during the development process can also lead to reliance on the hardware and software of that contractor, which may become impractical to sustain and difficult to transition to a new program when need arises. Proprietary software can bedevil future program development efforts.

Program development in this stage—whether for software or hardware—can follow several approaches, each of which has its own steps. Two prominent schools of thought are the “waterfall” and “agile” methodologies.²⁰⁶ Although we do not compare them in detail here, the key difference between the two is that agile development relies on iterations that deliver incremental capabilities, while waterfall development delivers a product at the end of the development cycle. Agile methodologies are increasingly popular within the Air Force and for AF DCGS applications, both within and beyond software development. The transition from waterfall to agile can also bring about dramatic reductions in timelines for the Air Force—cutting five-to-nine-year release cycles to a range of weeks or months. Additionally, agile development can blur the line between the development and launch stages of the implementation cycle because there is often iteration between the two (e.g., in a pilot program).

Stage 3: Launch the Program

Perhaps the most important point about this stage is that it exists at all: A new program that just shows up without warning will not be used,²⁰⁷ and, without active measures from leadership, it will not show up at all. A program's launch should be a planned process that appropriately conveys and disseminates information, trains all levels of users and supervisors, and rigorously evaluates progress toward clear goals.²⁰⁸ Successful launches often use a measured approach that

²⁰⁵ Robert H. Anderson, Tora K. Bikson, Rosalind Lewis, Joy S. Moini, and Susan G. Straus, *Effective Use of Information Technology: Lessons About State Governance Structures and Processes*, Santa Monica, Calif.: RAND Corporation, MR-1704-BSA, 2013.

²⁰⁶ Darrell K. Rigby, Jeff Sutherland, and Hirotaka Takeuchi, “Embracing Agile,” *Harvard Business Review*, May 2016.

²⁰⁷ We witnessed analysts use clever workarounds to avoid tools or workflows that did not meet their needs.

²⁰⁸ Laura J. Damschroder, Rosalind E. Keith, Susan R. Kirsh, Jeffery A. Alexander, and Julie C. Lowery, “Fostering Implementation of Health Services Research Findings into Practice: A Consolidated Framework for Advancing Implementation Science,” *Implementation Science*, Vol. 4, No. 50, August 2009; Tracy Ann Sykes and Jonathan L. Johnson, “Enterprise System Implementation and Employee Job Performance: Understanding the Role of Advice Networks,” *MIS Quarterly*, Vol. 38, No. 1, 2014.

incorporates dry runs, pilot studies, or a gradual introduction of the full process/program.²⁰⁹ Within the Air Force, this stage may take one to two years today (with individual pilot runs lasting approximately 30 days), although the specific details of an implementation plan will vary according to the program of interest and operational needs. Depending on the difficulty of the task, this could be shorter.

The importance of this preparation should not be underestimated. It is not enough simply to verify that a program functions according to its specifications. A lack of preparation or a belief that the new program's appeal is obvious and universal can both cause implementations to fail.²¹⁰ Once the program is advanced to beta testing or a pilot run, care must be taken to engage early users who are similar in background (professionally and personally) to the intended users.²¹¹ This is to help the implementation gain acceptance and check original assumptions regarding needs. If, as in the BlackBerry pilot (Figure 8.2), initial assumptions were incorrect and the pilot demonstrates that the program does not meet the needs, or if the pilot reveals that the program has a significant flaw requiring redevelopment or abandonment, then it will be less burdensome to learn these facts in a pilot test than after a full launch.

Figure 8.2. Case Study: Pilot BlackBerry Use for Law Enforcement

Stage 1: Law enforcement teams needed rapid access and communications capabilities to detect potential threats and coordinate actions.

Stage 2: Two law enforcement units decided to introduce BlackBerries among their squads.

Stage 3: Pilots studies at both sites revealed that some of the technical features of the BlackBerry were incompatible with the users' working conditions. For example, it was impractical and dangerous to input a password while pursuing a target. Neither site had training that was specific to law enforcement or that discussed organizational policies, resulting in confusion and misinformation.

The main difference between the two sites in enthusiasm for the BlackBerries had to do with a vision of the potential impact of the device propagated by an influential user-champion who led users in Site X to see themselves as becoming "wireless investigators of the 21st century" who can "direct all aspects of an [operation] from the field." This champion exhibited both charismatic and instrumental leadership..., secured resources, and guided a participatory implementation process...that revealed the benefits of the device, making its technical drawbacks more tolerable. Absent those factors, ergonomic difficulties and functionality limitations were much more salient in Site Y.^a

Managers at Site X mandated BlackBerry use and articulated a shared vision. Site Y had no champion and managers were ambivalent, taking a "wait and see" approach. Squads at Site Y with supervisors who encouraged BlackBerry use or other squad members who provided supplementary training were more enthusiastic about the device. The pilot study at Site X was more successful.

^a Susan G. Straus, Tora K. Bikson, Edward Balkovich, and John F. Pane, "Mobile Technology and Action Teams: Assessing Blackberry Use in Law Enforcement Units," *Computer Supported Cooperative Work (CSCW)*, Vol. 19, No. 1, 2009, pp. 45-71.

²⁰⁹ Damschroder et al., 2009.

²¹⁰ Leonard-Barton and Kraus, 1985.

²¹¹ Damschroder et al., 2009.

The success of a program's launch is highly influenced by several key players. The identities of these key players and the roles they fill will change and shift throughout the implementation process and the lifetime of the program. Exact definitions of these roles vary within the literature, but four types discussed here include *opinion leaders*, *champions*, *internal implementation leaders*, and *external change agents*.²¹² Note that an individual may fit more than one category. *Opinion leaders* can be technical experts who influence an implementation through their subject-matter expertise²¹³ or peers who are well-respected and liked within the organization.²¹⁴ *Champions* actively support the implementation, frequently serving to sway nonbelievers in the new program or process. Champions may be more successful at positively influencing an implementation if they possess some level of authority. *Internal implementation leaders* have formal responsibility for the implementation as part of their job duties, while *external change agents* are associated with an outside organization that may be assisting in some way with the implementation (such as representatives from a software company, management consultants, or contractors).²¹⁵ Importantly, external change agents hold a supporting and not a leading role.²¹⁶

Regardless of the roles they play, leadership must convey a strategic vision and demonstrate support for the new program. Buy-in of the vision by different stakeholder groups must also be achieved.²¹⁷ Effective buy-in for new AF DCGS analytic systems should not be assumed and cannot simply be commanded; this is an area where simply ordering change is insufficient. User engagement is essential.²¹⁸

Stage 4: Sustain the Program

For any system, investment must be sustained over time, well beyond the initial development or launch.²¹⁹ In addition to providing continued financial and personnel support, leadership must

²¹² Damschroder et al., 2009.

²¹³ Leonard-Barton, 1985.

²¹⁴ C. Tucker, "Identifying Formal and Informal Influence in Technology Adoption with Network Externalities," *Management Science*, Vol. 54, No. 12, August 2008, pp. 2024–2038.

²¹⁵ Jon (Sean) Jasperson, Pamela E. Carter, and Robert W. Zmud, "A Comprehensive Conceptualization of Post-Adoptive Behaviors Associated with Information Technology Enabled Work Systems," *MIS Quarterly*, Vol. 29, No. 3, September 2005; and Damschroder et al., 2009.

²¹⁶ External change agents can be effective. We heard several airmen speak appreciatively of "our" NRO representative. Such liaisons are surely one reason why NRO-led tools have been adopted by AF DCGS analysts.

²¹⁷ Kathrin M. Cresswell, David W. Bates, and Aziz Sheikh, David W Bates, and Aziz Sheikh, "Ten Key Considerations for the Successful Implementation and Adoption of Large-Scale Health Information Technology," *Journal of the American Medical Informatics Association*, Vol. 20, June 2013.

²¹⁸ The 480th ISRW was exemplary in this regard during the period of this study. We have seen more than ever before airmen we met across the enterprise who expressed that their wing's leadership was genuinely supportive of new ideas.

²¹⁹ Jasperson, Carter, and Zmud, 2005.

continue to reaffirm their support for the project by clearly informing (and reminding) users of the purpose and potential benefits of implementing a new program, as illustrated in the Health First case study (Figure 8.3).²²⁰ As the program implementation progresses, it must be evaluated periodically to determine whether the foreseen benefits have materialized as expected. The success of the program and the success of the program’s implementation—although related—are not synonymous and should both be evaluated.²²¹ Evaluation can be conducted through the tracking and analysis of relevant metrics, traditional feedback, and anecdotal or informal feedback.

Figure 8.3. Case Study: Health First

Stage 1: Health First, a hospital system in Florida, was suffering from inefficient practices in patient flow that led to miscommunication between hospitals and lost revenue.^a

Stage 2: They selected and implemented a hospital operations software that allowed them to track bed occupancy across hospitals in their system.

Stage 3: Health First hired a technical expert in Lean and Six Sigma to lead the implementation. External change agents from TeleTracking Technologies, Inc. and Central Patient Logistics provided support. Management articulated clear strategic goals to all staff and vocally supported the implementation over several months. Leadership emphasized continual process improvement, listened to feedback from staff, and created an environment to track and reward improved performance.

Software generated actionable, real-time data to spot issues & hold users accountable. Initial training was provided to teach staff how to use the new software. If employees were not using the program correctly or did not meet productivity goals, additional training was provided.

Stage 4: Health First noticed user resistance in the form of apathy toward new TeleTracking system. To address the resistance, managers persuaded the staff to embrace the new system by presenting data that described how the problem had previously affected the hospital and the benefits of the new changes.

The implementation and the program were both successful: Health First decreased patient wait times and improved care quality.

^a Blanchard, Janice C. and Robert S. Rudin, *Improving Hospital Efficiency Through Data-Driven Management: A Case Study of Health First, Florida, RR-1342-TELET*, Santa Monica, Calif.: RAND Corporation, 2015.

Example metrics for evaluation include initial acquisition speed, which is measured through the time it takes from initial contact with a commercial sector company until that company has a contract to partner with DoD.²²² Metrics can be evaluated both prior to rollout and throughout the software life cycle.²²³ The metrics should also try to capture user acceptance, which can be used

²²⁰ Leonard-Barton and Kraus, 1985.

²²¹ Damschroder et al., 2009.

²²² Defense Innovation Unit Experimental, *Annual Report 2017*, Silicon Valley, Calif.; Boston, Mass.; Austin, Tex.; and Washington, D.C., 2017.

²²³ Tajha Chappellet-Lanier, “Defense Innovation Board Proposes New Metrics for Assessing DOD Software Development,” *Fedscoop*, July 12, 2018.

to evaluate both program and implementation success and can be measured through feedback or usage metrics.²²⁴

Additionally, evaluations can inform users of how the program is affecting their output. If users can see the benefits for themselves, it can aid in building consensus and increasing commitment to the program.²²⁵ Conversely, evaluation may indicate that some users are struggling to adapt or are not adopting the program as intended. In this situation, supplementary or targeted training may be beneficial. Supervisors should also be trained in the new program their subordinates will be using.²²⁶ This not only allows the supervisors to better assist their employees but also helps prevent user resistance if the supervisors feel that their status is threatened by a perceived loss of control.²²⁷

Importantly, leadership should not expect a new program to pay immediate dividends; it takes time to learn a new system or process, and productivity often temporarily declines after introduction of a new technology.²²⁸ This expectation should be factored into any evaluation of the implementation and program performance. Analysis is a complex process with different interconnections; improvements in one area may not yield expected benefit without improvement in another, and some improvements are needed to lay the groundwork for AI/ML to come.

Stage 5: Abort When Necessary

No matter how smoothly the previous steps proceed, problems will almost certainly arise during the implementation period and throughout the lifetime of the program. Leadership will need to decide whether the program can be fixed or if it should come to an end. In addition to simple feedback, these issues are often brought to the attention of management through user resistance.²²⁹ Such resistance is not inherently positive or negative. Depending on how it manifests, resistance can allow users to convey concerns about the program or its effects, but it can also be destructive when it disrupts or prevents the adoption of a useful program. Resistance can take many forms, ranging from apathy to sabotage, and might occur because of perceived

²²⁴ Jasperson, Carter, and Zmud, 2005.

²²⁵ Leonard-Barton and Kraus, 1985.

²²⁶ Leonard-Barton and Kraus, 1985.

²²⁷ Liette Lapointe and Suzanne Rivard, "A Multilevel Model of Resistance to Information Technology Implementation," *MIS Quarterly*, Vol. 29, No. 3, September 2005.

²²⁸ Leonard-Barton and Kraus, 1985.

²²⁹ Rivard, Suzanne and Liette Lapointe, "Information Technology Implementers' Responses to User Resistance: Nature and Effects," *MIS Quarterly*, Vol. 36, No. 3, 2012.

threats of the program,²³⁰ mental inertia,²³¹ past experiences,²³² or personal characteristics.²³³ Many of the actions just discussed, such as identifying needs, defining goals, planning the implementation process, and providing training, can help prevent or reduce the impact of user resistance.

However, once resistance occurs, appropriate action by the implementers becomes crucial. Effective intervention can address problems to the satisfaction of users, leading to improvements in program productivity, while inappropriate responses can escalate destructive behavior and even result in program failure. Although user resistance can be problematic, it is important to remember that, at this stage, it is often a symptom of an underlying problem. Leaders must avoid the temptation to treat only the symptom and not the disease.

Common implementer responses to resistance include inaction, acknowledgment, rectification, and dissuasion.²³⁴ *Inaction* is defined as unawareness of the resistance or choosing not to act to acknowledge or address it. *Acknowledgment*, on the other hand, involves discussion of the resistance (such as by a task force or focus group) but is not followed by action to rectify the underlying causes that are discovered. In an analysis of more than 80 case studies, both approaches were ineffective and often led to increases in resistance. *Rectification*—where the response is intended to fix the issue or issues that are provoking resistance—can include redesigning the system, additional training, eliminating access to the old system, or developing a new system entirely. Rectification was effective in the case studies we analyzed only when the corrective action was aligned with the problem at hand. For example, rectification was effective when the resistance was caused by a confusing interface and the software was redesigned to be more user-friendly. For that same type of situation, when the response was to simply explain how to use the flawed system, the user resistance increased. This highlights the need for implementers to closely monitor feedback and understand the cause of resistance when it occurs. *Dissuasion* includes authoritative persuasion (reprimands and mandated use of the system), supportive persuasion (including reassurance, explanations, and benefit rationalization), and coercion (threatening users if the resistance is not stopped)—all of which were effective at decreasing resistance so long as the implementer was credible. When the implementers' response was not deemed credible because the implementers were untrustworthy, confusing, or simply reassured without taking action, user resistance increased.

²³⁰ Lapointe and Rivard, 2005.

²³¹ Sung S. Kim, "The Integrative Framework of Technology Use: An Extension and Test," *MIS Quarterly*, Vol. 33, No. 3, September 2009.

²³² Hee-Woong Kim and Atreyi Kankanhalli, "Investigating User Resistance to Information Systems Implementation: A Status Quo Bias Perspective," *MIS Quarterly*, Vol. 33, No. 3, September 2009.

²³³ Venkatesh et al., 2003.

²³⁴ Rivard and Lapointe, 2012.

The ideal response to user resistance varies according to the context of the situation, but credible dissuasion and rectification aligned to the problem at hand are both generally helpful in decreasing resistance. Note that some manifestations of rectification, such as developing a new system, may involve aborting the existing system, investing resources elsewhere, and cycling back to an early stage of the process. This kind of “positive failure” is often necessary and leads to better systems that are more aligned with the organizational need and context, particularly when the need has shifted.

Finally, we note that the life cycle of software is finite. Offboarding will be necessary. As civilian technology advances and airmen who are familiar with more-advanced systems in their private lives enter the AF DCGS, user resistance can be a sign that the useful life of a successful program is at its end or that the time is ripe for change.

Existing Air Force and Department of Defense Efforts

The stages of technology implementation discussed here are not intended to supplant existing structures, such as the formal Defense Acquisition Model.²³⁵ Such consideration goes beyond the scope of this project. Rather, the overview of the implementation life cycle is a framework to think through processes and activities that are already underway. The discussion of the individual stages suggests practices to improve performance and avoid pitfalls. The stages laid out are flexible enough to accommodate both traditional acquisition cycles and agile methods.²³⁶

However, to truly support technology implementation beyond a single project, the culture of the Air Force must change—and efforts are already underway to do so. The Defense Innovation Board, an independent advisory council to the Secretary of Defense that was founded in 2016, is working to shift DoD culture by integrating new perspectives. The board’s 16 draft recommendations on this topic include a recommendation to embrace a culture of experimentation.²³⁷ This involves simultaneously testing and assessing multiple approaches; incentivizing leadership to promote innovation and experimentation; and encouraging employees to be creative, vocal, and risk-tolerant. Other voices within the Air Force and PAF share this perspective and recommend achieving a culture of innovation by designing an information technology, acquisition, and security environment that allows AI/ML to thrive.²³⁸

Within the Air Force, multiple ongoing initiatives espouse these principles. Founded in 2017 with the mission to facilitate innovation among airmen, AFWERX lowers the barriers to

²³⁵ Department of Defense Instruction Number 5000.02, *Operation of the Defense Acquisition System*, January 7, 2015, incorporating change 3, August 10, 2017.

²³⁶ Suzanne Miller and Dan Ward, *Update 2016: Considerations for Using Agile in DoD Acquisition*, Pittsburgh, Pa.: Software Engineering Institute, Carnegie Mellon University, CMU/SEI-2016-TN-001, December 2016.

²³⁷ Defense Innovation Board, “Our Work,” webpage, undated.

²³⁸ Cortney Weinbaum and John N. T. Shanahan, “Intelligence in a Data-Driven Age,” *Joint Force Quarterly*, No. 90, third quarter, July 2018.

partnering with the Air Force by providing coaching to startups and pairing them directly with airmen to share early-stage feedback.²³⁹ AFWERX serves as a bridge between stovepipes of activity within the Air Force itself so that best practices can be shared. This is part of a larger series of WERX hubs that the Air Force is opening in conjunction with the Air Force Research Laboratory and others.²⁴⁰

The 480th ISRW has also taken steps to craft opportunities for its airmen to take assignments as “combat coders” or be embedded in the Silicon Valley ethos at Defense Innovation Unit Experimental (DIUx). The 480th ISRW has aimed to encourage incremental achievements (“roof shots” versus “moonshots”) and to have airmen-driven ideas generate conversation among senior leaders. Without this effort, 480th ISRW leadership felt that the rollout of new tools would amount to little more than “innovation theater.”²⁴¹

These efforts and others, such as Project Kessel Run (Figure 8.4), aim to foster a culture of innovation at the strategic and operational levels; they lay the groundwork for innovation. This is not an easy task—where other implementation obstacles discussed above can be alleviated with increased financial resources, cultural barriers are not surmounted as simply. Countering organizational inertia and driving change requires the repeated investment of time, strategic vision, and leadership support. Without this investment in a cultural shift to change underlying organizational assumptions and thinking, implementations often fail.

Figure 8.4. Case Study: Project Kessel Run

The Air Force's Kessel Run Experimentation Lab, set up in 2018 in Boston's North End, serves as an innovation hub for next generation combat software for the Air Operations Center (AOC). The lab consists of product developers divided into application teams, operations, infrastructure and support teams. Staff have 6-month temporary assignments to the lab. The lab is modeled after start-up commercial software development companies.

Kessel Run is a project run out of the Air Force Life Cycle Management Center to modernize the Air Operations Center, with DIUx's support, whereby over 70 airmen have recently undergone training through a partnership with a company, Pivotal Labs, to learn software and app development in a genuine Agile software development environment. It is DoD's version of a Software Factory. These airmen regularly ship new features every week in an iterative process seen in successful software companies. Kessel Run has already saved vast sums of money that would otherwise have been spent through the traditional acquisition process. Cycle times that may have extended years are accomplished in weeks.^a

^a Schmidt, 2018.

²³⁹ Samantha Ehlinger, “Air Force Innovation Group AFWERX Expands to Texas,” *Fedscoop*, June 29, 2018.

²⁴⁰ Associated Press, “Air Force Innovation Hub Launches in Alabama,” *Air Force Times*, September 2, 2018.

²⁴¹ Jason M. Brown, “Building an Innovation Ecosystem Part 1: Culture and Framework,” *Medium*, December 10, 2017.

Innovation Examples from the Intelligence Community

In addition to this general discussion, we collected several relevant examples of innovation within the IC. We visited sites to observe operations, interview analysts, and speak to SMEs at innovation centers or sites that have instituted innovation initiatives. To this end, the research team visited two NSA sites, five DGS sites, and the NRO and spoke to various staff at DIUx, Netflix, Google, and Air Force/A2. These cases provide specific examples of how some of the principles of innovation have been put into practice.

National Security Agency Hawaii

This facility supports several innovation-related initiatives that vary in scope and time. The Hix Incubation Cell, for example, applies a “shark tank” approach and awards short-timeline (under six months) projects using ideas submitted by resident staff. Part of the package reviewed by senior leaders includes staff skill mix and sufficiency of funding. An instructive feature of these proposals is that classification compliance is often “baked into” the project from the start, and progress reviews use “fail gates” to keep efforts on track and create off-ramps for aborting the effort if it does not appear that it will succeed in time. Part of the philosophy of these programs is that staff are to be rewarded for engaging rather than merely anticipate achieving the final outcome: The goal is to encourage risk-taking.

The organization also has MAD Scientist residents to apply data-science techniques in support of the mission. Finally, Code Junkies, a weekly forum, allows resident staff to learn new coding skills and helps accelerate workflow through quick-turn projects.

National Security Agency Georgia

This facility leverages a lab team of data scientists, programmers, and mathematicians who reside locally and provide direct support to the mission. In addition, DoD and NSA teams create an environment for leveraging IC databases, analytic tools, and information technology support. This colocation provides an unusual depth of expertise for this area.

National Reconnaissance Office

Although it is an acquisition organization for space-based national reconnaissance, NRO is structured for PED resource development of the IC and DoD. Resources exist to create an open architecture for PED capabilities and a tool-development environment that includes a research stream to explore future needs related to tasking, collection, and PED. Funding mechanisms are in place to leverage commercial small business efforts for quick-turn improvements. The organization is well-resourced to innovate in this area: The necessary processes, funding, and mechanisms to pursue improvements are in place.

Defense Innovation Unit Experimental

DoD established DIUx in 2015 to experiment with new ways to deliver innovative capabilities to the warfighter.²⁴² DIUx seeks to build relationships with leading technology firms and can also commit some of its own funds but does not seek to own what it helps create. In FY 2017, it had \$20 million to invest.²⁴³

One especially relevant initiative is Project Maven. This effort, which earlier involved Google,²⁴⁴ aims to leverage commercial ML advances to classify objects in imagery. The project involves data labeling, neural nets, compute (processing), program-of-record integration, and user engagement.²⁴⁵ The main thrust of the effort is to develop a learning data set by enlisting Air Force staff to tag more than a million visual images. Over the project's course, competitions have been held to encourage staff to participate in the tagging process. Project Maven aims to use an iterative approach to development that emphasizes user engagement on an ongoing basis.²⁴⁶

Data to Decisions

As of July 2018, relevant Data to Decisions (D2D) efforts were narrowly scoped to explore the utility of the D2D approach within the AF DCGS context. The approach to experimentation here is to conduct quarterly software sprints to automate processes identified by AF DCGS analysts as candidates for automation. Contractors (data scientists and programmers) are brought to a DGS site to quickly learn the problem, understand the work environment, and develop the proposed solution. A 2018 sprint at DGS-5 consisted of a short, two-week software sprint that needed further improvements downstream.

480th ISRW Innovation Labs

Within the past year, the 480th ISRW has opened innovation labs at various DGS sites. Some offer advanced hardware—at least one includes virtual headsets and a 3-D printer²⁴⁷—but the focus is on supporting coding activities. Access to these centers is generally open to everyone in the local ISRG. They provide 24/7 access for airmen interested in programming to learn, practice, and collaborate. To our knowledge, none have full-time staff but borrow from the local

²⁴² Carolyn Wong, “Enhancing ACC Collaboration with DIUx,” Santa Monica, Calif.: RAND Corporation, WR-1177-AF, 2017.

²⁴³ Wong, 2017.

²⁴⁴ Google has since pulled out this project (Wakabayashi and Shane, 2018).

²⁴⁵ Pellerin, 2017.

²⁴⁶ Gregory C. Allen, “Project Maven Brings AI to the Fight Against ISIS,” *Bulletin of the Atomic Scientists*, December 21, 2017.

²⁴⁷ Steve Hirsch, “Innovation Lab, with Star Trek Decal Opens at South Korea Base,” *Air Force Magazine*, undated.

ISRG. The open environment is conducive to creativity and, perhaps more important, indicates a concrete investment by leadership to help their human capital develop.

Successes

Several themes arose from the different innovation models within the IC and DoD innovation hubs. Leadership support for risk-taking was a consistent notable characteristic. As seen in the literature, fostering an environment that encourages taking risks and being creative starts at the top. User buy-in and an effective feedback loop to “fix” the initial version are also important elements. By working side-by-side with analysts, developers can better understand the environment, the problem, and how their tool fits into the workflow. We also saw that airmen may find coding rewarding to address immediate workflow problems, and some are interested in contributing in this way. We also noted that, with classification requirements and disparate systems, addressing compliance early in development will lower risk to schedule slippage and cost growth.

Conclusion

With appropriate care, planning, and support, the Air Force is poised to successfully implement technologies and enable AF DCGS analysts to take full advantage of their tools now and in the future. But how they implement these technologies is as important as what they implement. Taking care to plan for and implement each stage of the technology implementation cycle will help the AF DCGS succeed in this regard.

This overview of implementation is broad enough that it can be applied to new programs, processes, and tools of all scales. The discussion on a culture of innovation and continuous improvement is similarly adaptable, and local change at the airbase level can complement and enhance the impact of Air Force-wide cultural shifts. The guidance on creating a cultural shift complements the technology-implementation framework. Innovative managers support new initiatives and become implementation leaders. A culture that prizes continuous improvement trains airmen to generate new ideas and empowers them to see these ideas to fruition. The success of new programs and the retirement of outdated ones help create an Air Force that can adapt to new security realities.

Appendix A. Defining Technology Readiness Levels for Artificial Intelligence/Machine Learning

Technology readiness levels (TRLs) are widely used in government and industry to assess the maturity of technologies for development and procurement. First developed at NASA in the 1970s, TRLs were formalized into a seven-level scale in the 1980s and then expanded to the now-familiar nine-level scale in the early 1990s. DoD adapted this scale to inform its research and development and procurement processes and has continued to refine it since.²⁴⁸ Other organizations in both government and industry have also adapted TRLs for their purposes, but these are not always mutually compatible. For example, DoD defines *TRL 5* as “component and/or breadboard evaluation in a relevant environment,”²⁴⁹ while the oil industry imposes a stricter requirement at the same TRL: “full-scale prototype built and integrated into intended operating system with full interface and functionality tests.”²⁵⁰ It is also not always obvious how to apply this language, designed for hardware, to software embodying a mathematical construct such as an AI/ML algorithm.

TRLs are inherently specific to the technology being analyzed. The original NASA scale, for example, only addressed critical technology elements of space systems. Designing a TRL scale that can be applied informatively to a wide variety of technologies is challenging, and a common critique of the entire TRL methodology is that people attempt to apply a single TRL scale to an overly broad domain of technologies.²⁵¹ This is clearly an issue when attempting to apply the DoD scale to AI/ML technologies. For instance, it is not obvious what constitutes a “component and/or breadboard validation” for an ML algorithm. Moreover, it is sometimes unclear that an AI/ML application has even reached “TRL 1,” because algorithms can be specified that solve extremely difficult problems in theory but may not prove useful in practice.²⁵² The difficulty of predicting AI/ML development compounds these issues.

²⁴⁸ Mihály Héder, “From NASA to EU: The Evolution of the TRL Scale in Public Sector Innovation,” *The Innovation Journal: The Public Sector Innovation Journal*, Vol. 22, No. 2, 2017.

²⁴⁹ Assistant Secretary of Defense for Research and Engineering, *Technology Readiness Assessment (TRA) Guidance*, Washington, D.C.: U.S. Department of Defense, April 2011.

²⁵⁰ John Strutt and Don Wells, “API 17N—Recommended Practice for Subsea Production System Reliability, Technical Risk, and Integrity Management,” presented at the Offshore Technology Conference, Houston, Tex., May 5–8, 2014.

²⁵¹ Strutt and Wells, 2014.

²⁵² A particularly notorious example of this was Herbert Simon and Allen Newell’s General Problem Solver, which can theoretically solve an enormous class of well-specified problems but turned out to be almost useless in practice because it scales poorly to all but the smallest problem instances. See Peter Norvig, *Paradigms of Artificial Intelligence Programming: Case Studies in Common Lisp*, San Francisco, Calif.: Morgan Kaufmann, 1992, pp. 146–147.

Ideally, we would use a domain-specific TRL scale for AI/ML, but, to our knowledge, no such scale exists, and developing one would be beyond the scope of this project. Nevertheless, the standard DoD TRL scale can still be useful to draw analogies with the maturity of more-familiar technologies. In this spirit, we map the DoD TRL definitions to AI/ML as best we can (see Table A.1.). Note that this mapping refers to the algorithmic methods themselves and not to the specific applications based on them, which may be less mature. When we define technologies as being “mature” in this report, we are referring to those at levels 8 and 9 in this scale, (e.g., those that already have significant commercial deployment).

Table A.1. Proposed TRL Scale for AI/ML Methods

TRL	DoD Definition	AI/ML Interpretation	Example
0	—	Conjectural technology	AGI
1	Basic principles observed and reported	Proposed algorithm	Quantum neural network ^a
2	Technology concept and/or application formulated	Demonstration on toy problem or data set (e.g., Modified National Institute of Standards and Technology [MNIST ^b])	Neural Turing machine ^c
3	Analytical and experimental critical function and/or characteristic proof of concept	Demonstration on medium-scale problem or data set (e.g., Canadian Institute For Advanced Research [CIFAR-10 ^d], WordNet ^e)	Statistical-relational learning
4	Component and/or breadboard validation in a laboratory environment	Demonstration on significant problem or data set, such as ImageNet ^f	Reinforcement learning for robotics
5	Component and/or breadboard validation in a relevant environment	Demonstration on large problem or data set (e.g., VisualQA ^g)	Iterated-width planners
6	System or subsystem model or prototype demonstration in a relevant environment.	Test on “real-world” data set	Self-driving car
7	System prototype demonstration in an operational environment	Pilot commercial deployment	Deep learning for automated scene description
8	Actual system completed and qualified through test and demonstration	Significant commercial deployment	Deep learning for speech synthesis
9	Actual system proven through successful mission operations	Widespread commercial deployment	Neural machine language translation

^a An algorithm that could simulate a neural net on quantum computer hardware.

^b Yann LeCun, Corinna Cortes, and Christopher J. C. Burges, “The MNIST Database,” homepage, undated.

^c Mark Collier and Joeran Beel, “Implementing Neural Turing Machines,” in Věra Kůrková, Yannis Manolopoulos, Barbara Hammer, Lazaros Iliadis, and Ilias Maglogiannis, eds., *Artificial Neural Networks and Machine Learning—ICANN 2018, Proceedings of the 27th International Conference on Artificial Neural Networks, Rhodes, Greece, October 4–7, 2018, Part III*, Basel, Switzerland: Springer Nature Switzerland AG, 2018. ^b Alex Krizhevsky, *Learning Multiple Layers of Features from Tiny Images*, April 8, 2009.

^d Alex Krizhevsky, *Learning Multiple Layers of Features from Tiny Images*, April 8, 2009.

^e Christiane Fellbaum, ed., *WordNet: An Electronic Lexical Database*, Cambridge, Mass.: MIT Press, 1998.

^f Krizhevsky, Sutskever, and Hinton, 2012.

^g Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh, “Making the V in VQA Matter: Elevating the Role of Image Understanding in Visual Question Answering,” in *Proceedings of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

References

- 480th Intelligence, Surveillance and Reconnaissance Wing, “480th ISRW Institutes Combat Readiness Sustainment Program,” Air Combat Command webpage, October 30, 2017. As of July 23, 2020:
<https://www.acc.af.mil/News/Article-Display/Article/1376953/480th-isrw-institutes-combat-readiness-sustainment-program/>
- 548th Operational Support Squadron, *Sustainable DCGS: DGS-2 Pilot Study 18 Mar–25 Jun 2015*, Beale Air Force Base, Calif., July 2015.
- AI Impacts, “2016 Expert Survey on Progress in AI: Narrow Tasks,” webpage, undated. As of August 8, 2020:
https://aiimpacts.org/2016-expert-survey-on-progress-in-ai/#Narrow_tasks
- Air Combat Command Manual 14-401, *Air Force Distributed Common Ground System (DCGS) Training, Certification, and Quality Management*, Joint Base Langley-Eustis, Va., April 6, 2020. As of August 26, 2020:
<https://static.e-publishing.af.mil/production/1/acc/publication/accman14-401/accman14-401.pdf>
- Air Force Intelligence, Surveillance, and Reconnaissance Agency Instruction 14-153, *Air Force Distributed Common Ground System (AF DCGS) Operations Procedures*, March 15, 2013, 480th ISR Wing Supplement, February 5, 2014. As of September 3, 2018:
http://static.e-publishing.af.mil/production/1/480isrw/publication/afisrai14-153v3_480isrwsup_i/afisrai14-153v3_480isrwsup.pdf
- Alkire, Brien, Abbie Tingstad, Dale Benedetti, Amado Cordova, Irina Elena Danescu, William Fry, D. Scott George, Lawrence M. Hanser, Lance Menthe, Erik Nemeth, David Ochmanek, Julia Pollak, Jessie Riposo, Timothy William James Smith, and Alexander Stephenson, *Leveraging the Past to Prepare for the Future of Air Force Intelligence Analysis*, Santa Monica, Calif.: RAND Corporation, RR-1330-AF, 2016. As of July 23, 2020:
https://www.rand.org/pubs/research_reports/RR1330.html
- Allen, Gregory C., “Project Maven Brings AI to the Fight Against ISIS,” *Bulletin of the Atomic Scientists*, December 21, 2017. As of September 7, 2018:
<https://thebulletin.org/2017/12/project-maven-brings-ai-to-the-fight-against-isis/>
- Amann, Wayne, “Former AF ISR Agency Now Numbered Air Force,” U.S. Air Force webpage, 25th Air Force Public Affairs, October 2, 2014. As of August 20, 2018:

<https://www.af.mil/News/Article-Display/Article/503222/former-af-isr-agency-now-numbered-air-force/>

Anderson, Robert H., Tora K. Bikson, Rosalind Lewis, Joy S. Moini, and Susan G. Straus, *Effective Use of Information Technology: Lessons About State Governance Structures and Processes*, Santa Monica, Calif.: RAND Corporation, MR-1704-BSA, 2013. As of July 23, 2020:

https://www.rand.org/pubs/monograph_reports/MR1704.html

Arkin, William M., “Smart Bombs, Dumb Targeting?” *Bulletin of the Atomic Scientists*, Vol. 56, No. 3, May/June 2000, pp. 43–53.

Assistant Secretary of Defense for Research and Engineering, *Technology Readiness Assessment (TRA) Guidance*, Washington, D.C.: U.S. Department of Defense, April 2011. As of August 8, 2020:

<https://apps.dtic.mil/dtic/tr/fulltext/u2/a554900.pdf>

Associated Press, “Air Force Innovation Hub Launches in Alabama,” *Air Force Times*, September 2, 2018. As of September 3, 2018:

<https://www.airforcetimes.com/news/your-air-force/2018/09/02/air-force-innovation-hub-launches-in-alabama/>

Beale Air Force Base, “9th Intelligence Squadron,” webpage, May 22, 2012. As of August 29, 2018:

<https://www.beale.af.mil/Library/Fact-Sheets/Display/Article/279964/9th-intelligence-squadron/>

Beall, Abigail, “Visual Trick Fools AI into Thinking a Turtle Is Really a Rifle,” *New Scientist*, November 3, 2017. As of September 10, 2018:

<https://www.newscientist.com/article/2152331-visual-trick-fools-ai-into-thinking-a-turtle-is-really-a-rifle/>

Bengio, Yoshua, Aaron Courville, and Pascal Vincent, “Representation Learning: A Review and New Perspectives,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 35, No. 8, August 2013, pp. 1798–1828.

Birkler, John, *Untying Gulliver: Taking Risks to Acquire Novel Weapon Systems*, Santa Monica, Calif.: RAND Corporation, OP-268-OSD, 2009. As of July 23, 2020:

https://www.rand.org/pubs/occasional_papers/OP268.html

Blanchard, Janice C., and Robert S. Rudin, *Improving Hospital Efficiency Through Data-Driven Management: A Case Study of Health First, Florida*, Santa Monica, Calif.: RAND Corporation, RR-1342-TELET, 2015. As of September 3, 2018:

https://www.rand.org/pubs/research_reports/RR1342.html

- Bostrom, Nick, *Superintelligence: Paths, Dangers, Strategies*, New York: Oxford University Press, 2014.
- Brown, Jason M., “Operating the Distributed Common Ground System: A Look at the Human Factor in Net-Centric Operations,” *Air and Space Power Journal*, Winter 2009, pp. 51–57. As of August 21, 2018:
<http://www.dtic.mil/dtic/tr/fulltext/u2/a640256.pdf>
- Brown, Jason M., “Building an Innovation Ecosystem Part 1: Culture and Framework,” *Medium*, December 10, 2017. As of September 3, 2018:
<https://medium.com/@jasonmbro/building-an-innovation-ecosystem-part-1-culture-and-framework-95b275be4902>
- Brown, Jason M., and David Vernal, “Time-Dominant Fusion in a Complex World: Defining Time-Dominant Fusion and Its Interdependent Relationship with Airborne ISR Capabilities and Air Force DCGS,” *Trajectory*, November 11, 2014. As of August 23, 2018:
<http://trajectorymagazine.com/time-dominant-fusion-in-a-complex-world/>
- Butler, Amy, “Eyes Wide Open,” Aviation Week Network, September 19, 2011a, pp. 36–37.
- Butler, Amy, “USAF Turns to Hyperspectral Sensors in Afghanistan,” Aviation Week Network, September 19, 2011b. As of August 26, 2020:
<https://aviationweek.com/usaf-turns-hyperspectral-sensors-afghanistan>
- Capaccio, Tony, “Northrop Drone Flies Over Japan Reactor to Record Data,” *Bloomberg*, March 17, 2011. As of August 29, 2018:
<https://www.bloomberg.com/news/articles/2011-03-16/northrop-grumman-drone-to-fly-over-japan-reactor-to-gather-data>
- Carlini, Nicholas, and David Wagner, “Audio Adversarial Examples: Targeted Attacks on Speech-to-Text,” presented at the 39th IEEE Symposium on Security and Privacy, San Francisco, Calif., May 24, 2018.
- Chappellet-Lanier, Tajha, “Defense Innovation Board Proposes New Metrics for Assessing DOD Software Development,” *Fedscoop*, July 12, 2018. As of September 2, 2018:
<https://www.fedscoop.com/defense-innovation-board-software-development-metrics/>
- CNN, “Milosevic Accepts Peace Plan, Finnish Envoy Says,” June 3, 1999. As of August 29, 2018:
<http://www.cnn.com/WORLD/europe/9906/03/kosovo.peace.04/>
- Cohen, Rachel, “Gorgon Stare to Receive BLOS Upgrades While Air Force Explores Replacement,” *Inside Defense*, April 6, 2018. As of August 21, 2018:
<https://insidedefense.com/daily-news/gorgon-stare-receive-blos-upgrades-while-air-force-explores-replacement>

- Collier, Mark, and Joeran Beel, “Implementing Neural Turing Machines,” in Věra Kůrková, Yannis Manolopoulos, Barbara Hammer, Lazaros Iliadis, and Ilias Maglogiannis, eds., *Artificial Neural Networks and Machine Learning—ICANN 2018, Proceedings of the 27th International Conference on Artificial Neural Networks, Rhodes, Greece, October 4–7, 2018, Part III*, Basel, Switzerland: Springer Nature Switzerland AG, 2018, pp. 94–104.
- Cordova, Amado, Lindsay D. Millard, Lance Menthe, Robert A. Guffey, and Carl Rhodes, *Motion Imagery Processing and Exploitation (MIPE)*, Santa Monica, Calif.: RAND Corporation, RR-154-AF, 2013. As of July 23, 2020:
https://www.rand.org/pubs/research_reports/RR154.html
- Cresswell, Kathrin M., David W. Bates, and Aziz Sheikh, “Ten Key Considerations for the Successful Implementation and Adoption of Large-Scale Health Information Technology,” *Journal of the American Medical Informatics Association*, Vol. 20, June 2013, pp. e9–e13. As of September 2, 2018:
<https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3715363/>
- Damschroder, Laura J., David C. Aron, Rosalind E. Keith, Susan R. Kirsh, Jeffery A. Alexander, and Julie C. Lowery, “Fostering Implementation of Health Services Research Findings into Practice: A Consolidated Framework for Advancing Implementation Science,” *Implementation Science*, Vol. 4, No. 50, August 2009.
- Defense Advanced Research Projects Agency, *Broad Agency Announcement: Explainable Artificial Intelligence (XAI)*, Arlington, Va., DARPA-BAA-16-53, August 10, 2016. As of September 2, 2018:
<https://www.darpa.mil/program/explainable-artificial-intelligence>
- Defense Airborne Reconnaissance Office, *UAV Annual Report FY 1996*, Washington, D.C.: Office of the Under Secretary of Defense, Acquisition and Technology, November 6, 1996. As of September 14, 2018:
<http://www.dtic.mil/dtic/tr/fulltext/u2/a360512.pdf>
- Defense Innovation Board, “Our Work,” webpage, undated. As of August 8, 2020:
<https://innovation.defense.gov/Recommendations/>
- Defense Innovation Unit Experimental, *Annual Report 2017*, Silicon Valley, Calif.; Boston, Mass.; Austin, Tex.; and Washington, D.C., 2017. As of September 2, 2018:
https://assets.ctfassets.net/3nanhbfr0pc/61HMsvScYqxLsLIHtVWeAs/b7c29bce38233e502bef789a03643626/DIUx_2017_Annual_Report_FINAL.pdf
- Defense Science Board, *Lessons Learned During Operations Desert Shield and Desert Storm*, Washington, D.C., June 8, 1992.

- Department of Defense Instruction Number 5000.02, *Operation of the Defense Acquisition System*, January 7, 2015, incorporating change 3, August 10, 2017. As of September 2, 2018: <https://www.dau.mil/guidebooks/Shared%20Documents/DoDI%205000.02.pdf>
- Deptula, David A., keynote speech, C4ISR Journal Conference, Arlington, Va., October 2009.
- Deptula, David A., “Air Force ISR in a Changing World: Changing Paradigms While Optimizing ‘Low Density’ to Meet ‘High Demand,’” in Keith Brent, ed., *The Art of Air Power: Proceedings of the Royal Australian Air Force Air Power Conference*, Canberra, Australia: Commonwealth of Australia, March 30, 2010.
- Dickinson, David, “Navy Resumes Celestial Navigation Course,” *Sky and Telescope*, April 5, 2016. As of August 27, 2018: <https://www.skyandtelescope.com/astronomy-news/u-s-navy-resumes-celestial-navigation-training-04042016/>
- Didier, Jeremy, “Sustainable DCGS: DGS-2 Pilot Study 18 Mar – 25 Jun 2015,” 548th Operational Support Squadron, July 2015.
- DoD—See U.S. Department of Defense.
- Dodson, Brian, “DARPA’s New 1.8-Megapixel Camera is a Super High-Resolution Eye in the Sky,” *PBS*, February 11, 2013. As of August 7, 2020: http://arizonaenergy.org/News_13/News_Feb13/DARPAnew18-gigapixelcameraisasuperhigh-resolutioneyeinthesky.html
- Drezner, Jeffrey A., and Michael Simpson, *Exploring Parallel Development in the Context of Agile Acquisition: Analytical Support to the Air Superiority 2030 Enterprise Capability Collaboration Team*, Santa Monica, Calif.: RAND Corporation, RR-1808-AF, 2017. As of July 23, 2020: https://www.rand.org/pubs/research_reports/RR1808.html
- Ehlinger, Samantha, “Air Force Innovation Group AFWERX Expands to Texas,” *Fedscoop*, June 29, 2018. As of September 2, 2018: <https://www.fedscoop.com/afwerx-austin-texas-capital-factory-aetc-diux/>
- Elsayed, Gamaleldin F., Shreya Shankar, Brian Cheung, Nicolas Papernot, Alex Kurakin, Ian Goodfellow, and Jascha Sohl-Dickstein, “Adversarial Examples That Fool Both Computer Vision and Time-Limited Humans,” paper presented at the 32nd Conference on Neural Information Processing Systems, Montréal, Canada, December 4, 2018. As of July 23, 2020: <http://papers.nips.cc/paper/7647-adversarial-examples-that-fool-both-computer-vision-and-time-limited-humans>

- Environmental Systems Research Institute, “ArcGIS Pro 2.2 System Requirements,” webpage, undated. As of September 17, 2018:
<http://pro.arcgis.com/en/pro-app/get-started/arcgis-pro-system-requirements.htm>
- Environmental Systems Research Institute, “At the NGA, GIS Underpins Virtually Everything,” *ArcNews*, Spring 2017. As of September 4, 2018:
<http://www.esri.com/esri-news/arcnews/spring17/articles/at-the-nga-gis-underpins-virtually-everything>
- Feickert, Andrew, and Emma Chanlett-Avery, *Japan 2011 Earthquake: U.S. Department of Defense (DOD) Response*, Washington, D.C.: Congressional Research Service, R41690, June 2, 2011. As of July 23, 2020:
<https://fas.org/sgp/crs/row/R41690.pdf>
- Fellbaum, Christiane, ed., *WordNet: An Electronic Lexical Database*, Cambridge, Mass.: MIT Press, 1998.
- Foster, Stacy, “U-2 Reconnaissance Aircraft Deployed to Aid Japan Relief Efforts,” 51st Fighter Wing Public Affairs, U.S. Air Force, March 13, 2011. As of August 29, 2018:
<https://www.af.mil/News/Article-Display/Article/113963/u-2-reconnaissance-aircraft-deployed-to-aid-japan-relief-efforts/>
- General Atomics Aeronautical, “Lynx Multi-Mode Radar,” website, undated. As of September 17, 2018:
<http://www.ga-asi.com/lynx-multi-mode-radar>
- General Dynamics, “Multi-INT Analysis and Archive System (MAAS),” webpage, undated. As of September 3, 2018:
<https://gdmissionsystems.com/en/products/intelligence-systems/geospatial-intelligence/multi-int-analysis-and-archive-system>
- Giannetti, William, “A Commonsense Approach to Intelligence, Surveillance, and Reconnaissance Operations,” *Air and Space Power Journal*, Vol. 30, No. 3, Fall 2016, pp. 28–81.
- Gilmer, Justin, Ryan P. Adams, Ian Goodfellow, David Andersen, and George E. Dahl, “Motivating the Rules of the Game for Adversarial Example Research,” *arXiv preprint arXiv:1807.06732*, July 20, 2018.
- GitHub, “tesseract-ocr,” webpage, undated. As of August 9, 2020:
<https://github.com/tesseract-ocr/>
- Goodfellow, Ian, Jonathon Shlens, and Christian Szegedy, “Explaining and Harnessing Adversarial Examples,” paper presented at the International Conference on Learning Representations 2015, May 7–9, 2015, pp. 1–11.

- Goyal, Yash, Tejas Khot, Aishwarya Agrawal, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh, “Making the V in VQA Matter: Elevating the Role of Image Understanding in Visual Question Answering,” in *Proceedings of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. As of September 13, 2018:
<https://arxiv.org/pdf/1612.00837.pdf>
- Grace, Katja, John Salvatier, Allan Dafoe, Baobao Zhang, and Owain Evans, “Viewpoint: When Will AI Exceed Human Performance? Evidence from AI Experts,” *Journal of Artificial Intelligence Research*, Vol. 62, 2018, pp. 729–754.
- Grissom, Adam R., Caitlin Lee, and Karl P. Mueller, *Innovation in the United States Air Force: Evidence from Six Cases*, Santa Monica, Calif.: RAND Corporation, RR-1207-AF, 2016. As of September 2, 2018:
https://www.rand.org/pubs/research_reports/RR1207.html
- Hambling, David, “New Army Camera Promises Super-Wide Surveillance,” *Wired*, August 19, 2009. As of August 21, 2018:
<https://www.wired.com/2009/08/new-army-camera-promises-total-surveillance/>
- Harnett, Gavin S., Andrew J. Lohn, and Alexander P. Sedlack, “Adversarial Examples for Cost-Sensitive Classifiers,” working paper presented at the 33rd Conference on Neural Information Processing System (NeurIPS 2019), Vancouver, Canada, December 13, 2019. As of September 1, 2020:
<https://arxiv.org/pdf/1910.02095.pdf>
- Haugeland, John, *Artificial Intelligence: The Very Idea*, Cambridge, Mass.: MIT Press, 1985.
- Haugh, Timothy D., and Douglas W. Leonard, “Improving Outcomes: Intelligence, Surveillance, and Reconnaissance Assessment,” *Air and Space Power Journal*, Vol. 31, No. 4, Winter 2017. As of July 27, 2020:
https://www.airuniversity.af.mil/Portals/10/ASPJ/journals/Volume-31_Issue-4/SLP-Haugh_Leonard.pdf
- Héder, Mihály, “From NASA to EU: The Evolution of the TRL Scale in Public Sector Innovation,” *The Innovation Journal: The Public Sector Innovation Journal*, Vol. 22, No. 2, 2017.
- Hirsch, Steve, “Innovation Lab, with Star Trek Decal Opens at South Korea Base,” *Air Force Magazine*, undated. As of July 27, 2020:
<https://www.airforcemag.com/daily/hawaii-raptors-get-creative-dod-says-afghan-strategy-working-lockheed-bolsters-cyber-protections/>
- Hoffman, Michael, “Gorgon’s Gaze Set for Fall in Afghanistan,” *Air Force Times*, June 13, 2010.

- Hollock, Jennifer A., *Flexibility Versus Expertise: A Closer Look at the Employment of United States Air Force Imagery Analysts*, master's thesis, Maxwell Air Force Base, Ala.: Air Command and Staff College, Air University, October 2017. As of July 27, 2020: <https://apps.dtic.mil/dtic/tr/fulltext/u2/1047475.pdf>
- Holstein, Peter, "Airmen Resiliency Team Provides 480th ISR Wing with Medical, Psychological and Spiritual Care," Air Force Medical Service webpage, Surgeon General Office of Public Affairs, May 24, 2017. As of August 18, 2018: <https://www.airforcemedicine.af.mil/News/Display/Article/1192570/airmen-resiliency-team-provides-480th-isr-wing-with-medical-psychological-and-s/>
- Hornik, Kurt, "Approximation Capabilities of Multilayer Feedforward Networks," *Neural Networks*, Vol. 4, No. 2, 1991, pp. 251–257.
- Huang, Sandy, Nicolas Papernot, Ian Goodfellow, Yan Duan, and Pieter Abbeel, "Adversarial Attacks on Neural Network Policies," workshop paper presented at the Fifth International Conference on Learning Representations, Toulon, France, April 26, 2017. As of September 1, 2018: <https://arxiv.org/pdf/1702.02284.pdf>
- Hurlburt Field, "11th Special Operations Intelligence Squadron," webpage, March 28, 2017. As of August 21, 2018: <https://www.hurlburt.af.mil/About-Us/Fact-Sheets/Fact-Sheets/Article/495223/11th-intelligence-squadron/>
- Jamieson, Perry D., *Lucrative Targets: The U.S. Air Force in the Kuwaiti Theatre of Operations*, Washington, D.C.: Air Force History and Museums Program, 2001.
- Jaspersen, Jon (Sean), Pamela E. Carter, and Robert W. Zmud, "A Comprehensive Conceptualization of Post-Adoptive Behaviors Associated with Information Technology Enabled Work Systems," *MIS Quarterly*, Vol. 29, No. 3, September 2005, pp. 525–557.
- Joint Publication 2-01, *Joint and National Intelligence Support to Military Operations*, Washington, D.C., July 5, 2017.
- Joint Publication 3-60, *Joint Targeting*, Washington, D.C., January 31, 2013. As of August 8, 2020: https://www.justsecurity.org/wp-content/uploads/2015/06/Joint_Chiefs-Joint_Targeting_20130131.pdf
- Keaney, Thomas A., and Eliot A. Cohen, *Gulf War Air Power Survey: Summary Report*, Washington, D.C.: U.S. Government Printing Office, 1993.
- Kennedy, J. Michael, "U.S. Rushes Defenses to Israel: American Troops to Operate Two Patriot Batteries: Gulf War: The Move May Forestall an Immediate Retaliatory Strike by the Jewish

- State After Two Attacks on Tel Aviv. Allied Warplanes Step Up the Search for Scud Mobile Launchers,” *Los Angeles Times*, January 20, 1991. As of August 29, 2018:
http://articles.latimes.com/1991-01-20/news/mn-804_1_kuwaiti-patrol-boat-two-patriot-missile-batteries-iraqi-missile-attacks
- Kim, Hee-Woong, and Atreyi Kankanhalli, “Investigating User Resistance to Information Systems Implementation: A Status Quo Bias Perspective,” *MIS Quarterly*, Vol. 33, No. 3, September 2009, pp. 567–582.
- Kim, Sung S., “The Integrative Framework of Technology Use: An Extension and Test,” *MIS Quarterly*, Vol. 33, No. 3, September 2009, pp. 513–537.
- Kimminau, Jon, “A Culminating Point for Air Force Intelligence, Surveillance, and Reconnaissance,” *Air and Space Power Journal*, Vol. 26, No. 6, November–December 2012, pp. 113–129. As of October 22, 2020:
https://www.airuniversity.af.edu/Portals/10/ASPJ/journals/Volume-26_Issue-6/V-Kimminau.pdf
- Krause, Merrick E., and Jeffrey D. Smotherman, “An Interview with Assistant Secretary of Defense for Homeland Defense Paul McHale,” *Joint Forces Quarterly Forum*, No. 40, January 2006, pp. 10–15.
- Krizhevsky, Alex, *Learning Multiple Layers of Features from Tiny Images*, University of Toronto, April 8, 2009. As of September 13, 2018:
<https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>
- Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton, “ImageNet Classification with Deep Convolutional Neural Networks,” Fernando Pereira, Christopher J. C. Burges, Léon Bottou, and Kilian Q. Weinberger, eds., *Proceedings of the Advances in Neural Information Processing Systems 25*, 2012, pp. 1097–1105.
- Lambeth, Benjamin S., *NATO’s Air War for Kosovo: A Strategic and Operational Assessment*, Santa Monica, Calif.: RAND Corporation, MR-1365-AF, 2001. As of July 27, 2020:
https://www.rand.org/pubs/monograph_reports/MR1365.html
- Langley, John, *Occupational Burnout and Retention of Air Force Distributed Common Ground System (DCGS) Intelligence Personnel*, dissertation, Santa Monica, Calif.: Pardee RAND Graduate School, RGSD-306, 2012. As of July 27, 2020:
https://www.rand.org/pubs/rgs_dissertations/RGSD306.html
- Langley, Pat, “The Changing Science of Machine Learning,” *Machine Learning*, Vol. 82, No. 3, March 2011, pp. 275–279.
- Lapointe, Liette, and Suzanne Rivard, “A Multilevel Model of Resistance to Information Technology Implementation,” *MIS Quarterly*, Vol. 29, No. 3, September 2005, pp. 461–491.

- LeCun, Yann, Corinna Cortes, and Christopher J. C. Burges, “The MNIST Database,” homepage, undated. As of September 13, 2018:
<http://yann.lecun.com/exdb/mnist>
- leidos via PRNewswire via COMTEX, “SAIC Launches Advanced Intelligence Multimedia Exploitation Suite (AIMES),” press release, McLean, Va., November 1, 2010. As of September 3, 2018:
<https://investors.leidos.com/news-and-events/news-releases/press-release-details/2010/SAIC-Launches-Advanced-Intelligence-Multimedia-Exploitation-Suite-AIMES/default.aspx>
- Leonard-Barton, Dorothy, and William A. Kraus, “Implementing New Technology,” *Harvard Business Review*, November 1985. As of July 27, 2020:
<https://hbr.org/1985/11/implementing-new-technology>
- Lowenthal, Mark M., *Intelligence: From Secrets to Policy*, 5th ed., London: SAGE Press, 2012.
- Mandel, Mark David, Thomas Hone, and Sanford S. Terry, *Managing “Command and Control” in the Persian Gulf War*, Westport, Conn.: Praeger Publishers, 1996.
- Marangunic, Nikola, and Andrina Granic, “Technology Acceptance Model: A Literature Review from 1986 to 2013,” *Universal Access in the Information Society*, Vol. 14, No. 1, February 16, 2014, pp. 81–95.
- Maslach, Christina, Wilmar B. Schaufeli, and Michael P. Leiter, “Job Burnout,” *Annual Review of Psychology*, Vol. 52, February 2001, pp. 397–422.
- Menthe, Lance, Bart E. Bennett, Joel Kvitky, Elliot Axelband, George E. Hart, Michael Nixon, and Sherrill Lingel, *Enhancing the Processing, Exploitation, and Dissemination of Technical Electronic Intelligence*, Santa Monica, Calif.: RAND Corporation, 2015a, Not available to the general public.
- Menthe, Lance, Amado Cordova, Elliot Axelband, Lindsay D. Millard, Abbie Tingstad, Endy M. Daehner, Kirsten M. Keller, and John Langley, *Technologies and Processes for Automating Processing, Exploitation, and Dissemination*, Santa Monica, Calif.: RAND Corporation, 2015b, Not available to the general public.
- Menthe, Lance, Amado Cordova, Carl Rhodes, Rachel Costello, and Jeffrey Sullivan, *The Future of Air Force Motion Imagery Exploitation: Lessons from the Commercial World*, Santa Monica, Calif.: RAND Corporation, TR-1133-AF, 2012. As of July 27, 2020:
https://www.rand.org/pubs/technical_reports/TR1133.html
- Menthe, Lance, Dahlia Anne Goldfeld, Abbie Tingstad, Sherrill Lingel, Edward Geist, Donald Brunk, Amanda Wicker, Sarah Soliman, Balys Gintautas, Anne Stickells, and Amado Cordova, *Technology Innovation and the Future of Air Force Intelligence Analysis: Volume*

- 1, *Findings and Recommendations*, Santa Monica, Calif.: RAND Corporation, RR-A341-1, 2021.
- Menthe, Lance, Dahlia Anne Goldfeld, Sherrill Lingel, Abbie Tingstad, and Anne Stickells, *Technology Innovation and the Future of Air Force Intelligence Analysis: Volume 3, Technical Assessment of Data Flow Maps*, Santa Monica, Calif.: RAND Corporation, forthcoming, Not available to the general public.
- Messitte, Nick, “How Avid Hopes to Fix a Broken Music Industry,” *Forbes*, April 30, 2015. As of September 3, 2018:
<https://www.forbes.com/sites/nickmessitte/2015/04/30/how-avid-hopes-to-fix-a-broken-music-industry/#17f1675d2a25>
- Metz, Cade, “Google Makes Its Special A.I. Chips Available to Others,” *New York Times*, February 12, 2018. As of January 31, 2020:
<https://www.nytimes.com/2018/02/12/technology/google-artificial-intelligence-chips.html>
- Miller, Suzanne, and Dan Ward, *Update 2016: Considerations for Using Agile in DoD Acquisition*, Pittsburgh, Pa.: Software Engineering Institute, Carnegie Mellon University, CMU/SEI-2016-TN-001, December 2016. As of September 2, 2018:
https://resources.sei.cmu.edu/asset_files/TechnicalNote/2016_004_001_484651.pdf
- Minsky, Marvin, ed., *Semantic Information Processing*, Cambridge, Mass.: MIT Press, 1968.
- Minsky, Marvin, and Seymour A. Papert, *Perceptrons: An Introduction to Computational Geometry*, Cambridge, Mass.: MIT Press, January 1969.
- mIRC.com, “Personal FAQ,” webpage, 2020. As of August 7, 2020:
<https://www.mirc.com/pfaq.html>
- Mnih, Volodymyr, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dhharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis, “Human-Level Control Through Deep Reinforcement Learning,” *Nature*, Vol. 518, No. 7540, February 25, 2015, pp. 529–533.
- Montes, Alexandre, “AF NTI Training Streamlines Intel Airmen to Mission,” press release, 70th ISRW Public Affairs, Fort George G. Meade, Md., November 21, 2016. As of July 27, 2020:
<https://www.nellis.af.mil/News/Article/1010459/af-nti-training-streamlines-intel-airmen-to-mission/>
- Nakashima, Ellen, and Craig Whitlock, “With Air Force’s Gorgon Drone ‘We Can See Everything,’” *Washington Post*, January 2, 2011. As of June 13, 2018:
<http://www.washingtonpost.com/wp-dyn/content/article/2011/01/01/AR2011010102690.html>

- National Air and Space Intelligence Center, “About Us: National Air and Space Intelligence Center,” webpage, May 2018. As of September 19, 2018:
<https://www.nasic.af.mil/About-Us/>
- National System for Geospatial-Intelligence, *Geospatial Intelligence (GEOINT) Basic Doctrine*, Publication 1.0, Springfield, Va., April 2018.
- NATO—See North Atlantic Treaty Organization.
- North Atlantic Treaty Organization, “The Aims of the Air Campaign,” October 30, 2000. As of August 29, 2018:
<https://www.nato.int/kosovo/repo2000/aims.htm>
- Norvig, Peter, *Paradigms of Artificial Intelligence Programming: Case Studies in Common Lisp*, San Francisco, Calif.: Morgan Kaufmann, 1992.
- Ochmanek, David, “Promoting Innovation and Modernization Within the Air Force,” Santa Monica, Calif.: RAND Corporation, RB-99-AF, 2003. As of September 2, 2018:
https://www.rand.org/pubs/research_briefs/RB99/index1.html
- Office of the National Geospatial-Intelligence Agency Historian, *The Advent of the National Geospatial-Intelligence Agency*, Washington, D.C., September 2011.
- Ostrowski, Kris Anthony, *Psychological Health Outcomes Within USAF Remotely Piloted Aircraft Support Career Fields*, dissertation, Daytona Beach, Fla.: Embry-Riddle Aeronautical University, June 2016. As of April 20, 2018:
<https://commons.erau.edu/cgi/viewcontent.cgi?article=1202&context=edt>
- Parloff, Roger, “Why Deep Learning Is Suddenly Changing Your Life,” *Fortune*, September 28, 2016. As of July 27, 2020:
<https://fortune.com/longform/ai-artificial-intelligence-deep-machine-learning/>
- Patrizio, Andy, “SSDs Get Bigger, While Prices Get Smaller,” *Network World*, May 22, 2018. As of September 1, 2018:
<https://www.networkworld.com/article/3274655/storage/ssds-get-bigger-while-prices-get-smaller.html>
- Pellerin, Cheryl, “Project Maven to Deploy Computer Algorithms to War Zone by Year’s End,” press release, Washington, D.C.: U.S. Department of Defense, July 21, 2017. As of July 3, 2018:
<https://www.defense.gov/News/Article/Article/1254719/project-maven-to-deploy-computer-algorithms-to-war-zone-by-years-end/>
- Pomerleau, Mark, “Carlisle: Overworked Airmen Can’t Train for Future Threats,” *Defense Systems*, September 18, 2015. As of August 29, 2018:

<https://defensesystems.com/Articles/2015/09/18/Hawk-Carlisle-Air-Force-training-shortfall.aspx>

Prince, Lillian, Wayne L. Chappelle, Kent D. McDonald, Tanya Goodman, Sara Cowper, and William Thompson, “Reassessment of Psychological Distress and Post-Traumatic Stress Disorder in United States Air Force Distributed Common Ground System Operators,” *Military Medicine*, Vol. 180, No. 3 Supplement, March 2015, pp. 171–178.

Rigby, Darrell K., Jeff Sutherland, and Hirotaka Takeuchi, “Embracing Agile,” *Harvard Business Review*, May 2016. As of September 2, 2018:
<https://hbr.org/2016/05/embracing-agile>

Rivard, Suzanne, and Liette Lapointe, “Information Technology Implementers’ Responses to User Resistance: Nature and Effects,” *MIS Quarterly*, Vol. 36, No. 3, 2012, pp. 897–920.

Robson, Seth, “Global Hawk Invaluable After Japan Disasters,” *Stars and Stripes*, September 12, 2011. As of July 27, 2020:
<https://www.stripes.com/news/global-hawk-invaluable-after-japan-disasters-1.154890>

Role, Lynette M., “New Artificial Intelligence Technology Assists Air Commandos with Decision-Making,” press release, Hurlburt Field, Fla.: Air Force Special Operations Command Public Affairs, September 13, 2017. As of July 3, 2018:
<http://www.afsoc.af.mil/News/Article-Display/Article/1309844/new-artificial-intelligence-technology-assists-air-commandos-with-decision-maki/>

Rosenau, William, *Special Operations Forces and Elusive Enemy Ground Targets: Lessons from Vietnam and the Persian Gulf War*, Santa Monica, Calif.: RAND Corporation, MR-1408-AF, 2001. As of July 27, 2020:
https://www.rand.org/pubs/monograph_reports/MR1408.html

Roser, Max, and Hannah Ritchie, “Technological Progress,” Our World in Data webpage, 2013. As of September 1, 2018:
<https://ourworldindata.org/technological-progress>

Rumelhart, David E., Geoffrey E. Hinton, and Ronald J. Williams, “Learning Representations by Back-Propagating Errors,” *Nature*, Vol. 323, October 9, 1986, pp. 533–536.

Russell, Stuart, and Peter Norvig, *Artificial Intelligence: A Modern Approach*, 2d ed., Upper Saddle River, N.J.: Prentice Hall, 2002.

Samuel, Arthur, “Some Studies in Machine Learning Using the Game of Checkers,” *IBM Journal of Research and Development*, Vol. 3, No. 3, July 1959, pp. 210–229.

Schmidt, Eric, “Statement of Dr. Eric Schmidt, House Armed Services Committee,” Washington, D.C.: U.S. House of Representatives, April 17, 2018. As of July 27, 2020:

<https://docs.house.gov/meetings/AS/AS00/20180417/108132/HHRG-115-AS00-Wstate-SchmidtE-20180417.pdf>

Schmitt, Eric, “Weak Serb Defense Puzzles NATO,” *New York Times*, March 26, 1999. As of August 8, 2020:

<https://archive.nytimes.com/www.nytimes.com/library/world/europe/032699kosovo-assess.html>

Sellber, Will, “The Other Side of the COIN,” *Air and Space Power Journal*, Vol. 32, No. 3, Fall 2018, pp. 72–84. As of February 8, 2019:

https://www.airuniversity.af.edu/Portals/10/ASPJ/journals/Volume-32_Issue-3/V-Selber.pdf

Siegelmann, Hava T., and Eduardo D. Sontag, “On the Computational Power of Neural Nets,” in *Proceedings of the Fifth Annual Workshop on Computational Learning Theory*, Pittsburgh, Pa., 1992, pp. 440–449.

Silver, David, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis, “Mastering the Game of Go with Deep Neural Networks and Tree Search,” *Nature*, Vol. 529, No. 7587, January 27, 2016, pp. 484–489.

Silver, David, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dhharshan Kumaran, Thore Graepel, Timothy Lillicrap, Karen Simonyan, and Demis Hassabis, “Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm,” *arXiv preprint arXiv:1712.01815*, 2017.

Sixteenth Air Force (Air Forces Cyber), “480th ISR Wing,” webpage, September 8, 2019. As of July 27, 2020:

<https://www.16af.af.mil/About-Us/Fact-Sheets/Display/Article/1963035/480th-isr-wing/>

Standage, Tom, “Taking Flight,” *The Economist: Technology Quarterly*, June 8, 2017. As of September 2, 2018:

<https://www.economist.com/technology-quarterly/2017/06/10/commercial-drones-are-the-fastest-growing-part-of-the-market>

Straus, Susan G., Tora K. Bikson, Edward Balkovich, and John F. Pane, “Mobile Technology and Action Teams: Assessing Blackberry Use in Law Enforcement Units,” *Computer Supported Cooperative Work*, Vol. 19, No. 1, 2010, pp. 45–71.

Strutt, John, and Don Wells, “API 17N—Recommended Practise for Subsea Production System Reliability, Technical Risk, and Integrity Management,” presented at the Offshore Technology Conference, Houston, Tex., May 5–8, 2014.

- Sykes, Tracy Ann, and Jonathan L. Johnson, "Enterprise System Implementation and Employee Job Performance: Understanding the Role of Advice Networks," *MIS Quarterly*, Vol. 38, No. 1, 2014, pp. 51–72.
- Tadjdeh, Yasmin, "Algorithmic Warfare: Google Versus the Pentagon, the Fallout," *National Defense Magazine*, August 2, 2018. As of August 21, 2018:
<https://www.nationaldefensemagazine.org/articles/2018/8/2/google-versus-the-pentagon-the-fallout>
- Thompson, Loren, "Air Force's Secret 'Gorgon Stare' Program Leaves Terrorists Nowhere to Hide," *Forbes*, April 10, 2015. As of August 24, 2018:
<https://www.forbes.com/sites/lorenthompson/2015/04/10/air-forces-secret-gorgon-stare-program-leaves-terrorists-nowhere-to-hide/#5cec252e7be4>
- Tingstad, Abbie, Dahlia Anne Goldfeld, Lance Menche, Robert A. Guffey, Zachary Haldeman, Krista S. Langeland, Amado Cordova, Elizabeth M. Waina, and Balys Gintautas, *Assessing the Value of Intelligence, Collected by U.S. Air Force Airborne Intelligence, Surveillance, and Reconnaissance Platforms*, Santa Monica, Calif.: RAND Corporation, RR-2742-AF, forthcoming.
- Tirpak, John A., "'PED Is Dead': ISR Roadmap Reaches Long for New Tech," *Air Force Magazine*, August 2, 2018. As of August 23, 2018:
<http://airforcemag.com/Features/Pages/2018/August%202018/PED-is-Dead-ISR-Roadmap-Reaches-Long-for-New-Tech.aspx>
- Torgeson, Igor, "Editing Like an Oscar Winner: Why Learn Avid Media Composer?" blog post, New York Film Academy, March 5, 2018. As of September 3, 2018:
<https://www.nyfa.edu/student-resources/editing-like-oscar-winner-learn-avid-media-composer-2/>
- Tucker, C., "Identifying Formal and Informal Influence in Technology Adoption with Network Externalities," *Management Science*, Vol. 54, No. 12, August 2008, pp. 2024–2038.
- U.S. Air Force, "Air Force Support to Hurricane Katrina/Rita Relief Operations: By the Numbers," Washington, D.C.: Headquarters U.S. Air Force, October 2005.
- U.S. Air Force, "U-2S/TU-2S," webpage, September 23, 2015a. As of August 21, 2018:
<https://www.af.mil/About-Us/Fact-Sheets/Display/Article/104560/u-2stu-2s/>
- U.S. Air Force, "Air Force Distributed Common Ground System," webpage, October 13, 2015b. As of August 16, 2018:
<https://www.af.mil/About-Us/Fact-Sheets/Display/Article/104525/air-force-distributed-common-ground-system/>

- U.S. Air Force, *Air Force Specialty Code 1N0X1: All Source Intelligence Analyst Career Field Education and Training Plan*, Washington, D.C.: Department of the Air Force, September 26, 2016. As of August 7, 2020:
<https://pdf4pro.com/view/department-of-the-air-force-cfotp-1n0x1-47b6d6.html>
- U.S. Department of Defense, *Conduct of the Persian Gulf War Final Report to Congress*, Washington, D.C.: U.S. Government Printing Office, 1992.
- U.S. Joint Chiefs of Staff, Joint Doctrine Division, *DOD Dictionary of Military and Associated Terms*, Washington, D.C., June 2018.
- Venkatesh, Viswanath, Michael G. Morris, Gordon B. Davis, and Fred D. Davis, “User Acceptance of Information Technology: Toward a Unified View,” *MIS Quarterly*, Vol. 27, No. 3, September 2003, pp. 425–478.
- Wakabayashi, Daisuke, and Scott Shane, “Google Will Not Renew Pentagon Contract That Upset Employees,” *New York Times*, June 1, 2018. As of July 3, 2018:
<https://www.nytimes.com/2018/06/01/technology/google-pentagon-project-maven.html>
- Weinbaum, Courtney, and John N. T. Shanahan, “Intelligence in a Data-Driven Age,” *Joint Force Quarterly*, No. 90, third quarter, July 2018, pp. 4–9.
- Weisgerber, Marcus, “The Pentagon’s New Artificial Intelligence Is Already Hunting Terrorists,” *Defense One*, December 21, 2017. As of August 25, 2018:
<https://www.defenseone.com/technology/2017/12/pentagons-new-artificial-intelligence-already-hunting-terrorists/144742/>
- Williams, Heather J., and Ilana Blum, *Defining Second Generation Open Source Intelligence (OSINT) for the Defense Enterprise*, Santa Monica, Calif.: RAND Corporation, RR-1964-OSD, 2018. As of July 27, 2020:
https://www.rand.org/pubs/research_reports/RR1964.html
- Williams, Sam, *Arguing A.I.: The Battle for Twenty-First-Century Science*, New York: AtRandom, 2002.
- Wong, Carolyn, “Enhancing ACC Collaboration with DIUx,” Santa Monica, Calif.: RAND Corporation, WR-1177-AF, 2017. As of September 3, 2018:
https://www.rand.org/pubs/working_papers/WR1177.html
- Yampolskiy, Roman V., “Turing Test as a Defining Feature of AI-Completeness,” in Xin-She Yang, ed., *Artificial Intelligence, Evolutionary Computing and Metaheuristics*, Berlin: Springer-Verlag, 2012, pp. 3–17.
- Zoph, Barret, and Quoc V. Le, “Neural Architecture Search with Reinforcement Learning,” talk presented at the 2017 International Conference on Learning Representations, Toulon, France, 2017.



Intelligence collections and demand have grown over the past two decades, and intelligence analysts are often performing routine tasks, leaving them unable to conduct larger strategic analyses that are needed to address future threats as outlined by the 2018 National Defense Strategy. The authors provide an in-depth analysis of technologies that could help the Air Force Distributed Common Ground System (AF DCGS) become more effective, efficient, adept at using human capital, and agile. A key point is that artificial intelligence (AI) and machine learning (ML) technologies alone do not solve these intelligence challenges; rather, if they are properly implemented and complemented by human analysts who have the right skills and training, the capabilities can allow the AF DCGS to evolve to better meet warfighter needs.

This is the second volume in a series about how AI/ML technology can help the AF DCGS meet the challenges of a demanding intelligence environment and the complexity of future threats envisioned by the 2018 National Defense Strategy. The authors provide more in-depth discussion of project methodology; a primer on AI/ML technology; case studies of analytic challenges in previous operations; best practices for successfully deploying new technologies; and other topics of interest to specialists, stakeholders, and experts.

\$35.00

ISBN-10 1-9774-0633-5
ISBN-13 978-1-9774-0633-0



www.rand.org